

EV355228351

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Network Load Balancing
with Connection Manipulation**

Inventor(s):

Abolade Gbadegesin

Sean B. House

Aamer Hydrie

Joseph M. Joy

Sanjay N. Kaniyar

Robert V. Welland

ATTORNEY'S DOCKET NO. **MS1-1517US**

1

Network Load Balancing with Connection Manipulation

2

3

4

5

RELATED PATENT APPLICATIONS

6 This U.S. Nonprovisional Application for Letters Patent (i) is a
7 continuation-in-part of co-pending U.S. Nonprovisional Application for Letters
8 Patent No. 10/610,506 (filed June 30, 2003), (ii) is a continuation-in-part of co-
9 pending U.S. Nonprovisional Application for Letters Patent No. 10/610,519 (filed
10 June 30, 2003), and (iii) is a continuation-in-part of co-pending U.S.
11 Nonprovisional Application for Letters Patent No. 10/610,321 (filed June 30,
12 2003).

13 Specifically, this U.S. Nonprovisional Application for Letters Patent is a
14 continuation-in-part of, and hereby incorporates by reference herein the entire
15 disclosure of, co-pending U.S. Nonprovisional Application for Letters Patent No.
16 10/610,506, filed June 30, 2003, and entitled “Flexible Network Load Balancing”.

17 Specifically, this U.S. Nonprovisional Application for Letters Patent is also
18 a continuation-in-part of, and hereby incorporates by reference herein the entire
19 disclosure of, co-pending U.S. Nonprovisional Application for Letters Patent No.
20 10/610,519, filed June 30, 2003, and entitled “Network Load Balancing with Host
21 Status Information”.

22 Specifically, this U.S. Nonprovisional Application for Letters Patent is also
23 a continuation-in-part of, and hereby incorporates by reference herein the entire
24 disclosure of, co-pending U.S. Nonprovisional Application for Letters Patent No.

1 10/610,321, filed June 30, 2003, and entitled "Network Load Balancing with
2 Session Information".

3

4 **TECHNICAL FIELD**

5 This disclosure relates in general to network load balancing and in
6 particular, by way of example but not limitation, to network load balancing with
7 connection manipulation, such as connection migration with tunneling and/or
8 connection migration in conjunction with application-level load balancing.

9

10 **BACKGROUND**

11 Communication, and many facets of life that involve communication, has
12 been greatly impacted by the Internet. The Internet enables information to be
13 communicated between two people and/or entities quickly and relatively easily.
14 The Internet includes many network nodes that are linked together such that
15 information may be transferred between and among them. Some network nodes
16 may be routers that propagate a packet from one link to another, may be individual
17 client computers, may be personal networks for different entities (e.g., intranets
18 for businesses), and so forth.

19 For this personal network case, as well as others, packets arriving at an
20 Internet node or nodes are distributed to other nodes within the personal network.
21 Such a personal network may be formed, for example, from a set of servers that
22 can each work on packets that arrive at the personal network. A business, a
23 university, a government office, etc. may receive many packets in a short
24 timeframe at its personal network. In order to respond in a timely manner and to
25 reduce the likelihood of rejection or loss of arriving packets, the personal network

1 may rely on multiple servers that can each work on the arriving packets
2 simultaneously.

3 The arriving packets are often inquiries pertaining to certain information,
4 such as a document, a catalog item, a web page, and so forth. The arriving packets
5 can also pertain to an economic transaction between a customer and a merchant.
6 Other purposes for the packets of a packet-based communication are possible.
7 Regardless, the arriving packets are distributed among different servers of a set of
8 servers to accommodate a rapid arrival of the packets and/or complex
9 communication exchanges.

10 The distribution of arriving packets among different servers of a set of
11 servers is often termed network load balancing. In other words, a load balancing
12 operation may be performed on packets as they arrive at a node or nodes of the
13 Internet when the node or nodes constitute a personal network and/or when they
14 connect the personal network to the Internet.

15 Such a load balancing operation is accomplished using dedicated hardware
16 that fronts the personal network at the node or nodes that connect the personal
17 network to the Internet and/or that provide a presence for the personal network on
18 the Internet. The physical hardware that performs the load balancing operation is
19 usually duplicated in its entirety to realize redundancy and improve availability of
20 the load balancing operation. To increase capacity for load balancing operations,
21 more-powerful hardware that replicates the entirety of the previous load balancing
22 hardware, and thus the operational capability thereof, is substituted for the
23 previous load balancing hardware. Such scaling up of the load balancing
24 operational capabilities is therefore confined to increasing the power of the
25 hardware via substitution thereof.

1 To implement a load balancing operation, the hardware usually performs a
2 round robin distribution of arriving connection requests. In other words, arriving
3 connection requests are distributed to servers of a set of servers in a linear,
4 repeating manner with a single connection request being distributed to each server.
5 This round-robin load balancing distribution of connections is typically utilized
6 irrespective of the condition of the personal network or the nature of an arriving
7 connection request. If a load balancing operation does extend beyond a round
8 robin distribution, these other factors are only considered to the extent that they
9 may be inferred from network traffic and/or from a congestion level of the
10 personal network.

11 Accordingly, there is a need for schemes and/or techniques that improve
12 network load balancing and/or the options associated therewith.

13

14 **SUMMARY**

15 In a first exemplary device implementation, a device includes: a
16 connection migrator that is configured to migrate connections away from the
17 device; the connection migrator capable of precipitating a compilation of protocol
18 state for a connection across a protocol stack; the connection migrator adapted to
19 aggregate the compiled protocol state with data for the connection into an
20 aggregated connection state of the connection; the connection migrator further
21 capable of causing the aggregated connection state to be sent toward a target
22 device.

23 In a first exemplary media implementation, one or more processor-
24 accessible media include processor-executable instructions that, when executed,
25 direct a device to perform actions including: receiving a connection state for a

1 connection; injecting the connection state for the connection into a network stack;
2 and continuing the connection using the injected connection state.

3 In a second exemplary media implementation, one or more processor-
4 accessible media include processor-executable instructions that, when executed,
5 direct a device to perform actions including: obtaining at least a portion of a
6 source/destination pair from an incoming packet; accessing an encapsulation
7 mapping table using the obtained at least a portion of the source/destination pair to
8 locate an encapsulation mapping entry; extracting a flow identifier from the
9 located encapsulation mapping entry; and replacing part of the incoming packet
10 with the extracted flow identifier to produce an encapsulated packet.

11 In a second exemplary device implementation, a device includes: a
12 tunneler that is configured to tunnel packets into the device; the tunneler having
13 access to an encapsulation mapping table, the encapsulation mapping table
14 including a plurality of encapsulation mapping entries, each encapsulation
15 mapping entry linking a flow identifier to at least a portion of a source/destination
16 pair; the tunneler adapted to accept an encapsulated packet having a particular
17 flow identifier; the tunneler capable of looking up a particular source/destination
18 pair at a particular encapsulation mapping entry using the particular flow
19 identifier; wherein the tunneler is further adapted to de-encapsulate the
20 encapsulated packet by replacing the particular flow identifier with at least part of
21 the particular source/destination pair.

22 Other method, system, approach, apparatus, application programming
23 interface (API), device, media, procedure, arrangement, etc. implementations are
24 described herein.
25

BRIEF DESCRIPTION OF THE DRAWINGS

The same numbers are used throughout the drawings to reference like and/or corresponding aspects, features, and components.

FIG. 1 is an exemplary network load balancing paradigm that illustrates a load balancing infrastructure and multiple hosts.

FIG. 2 is an exemplary network load balancing paradigm that illustrates multiple load balancing units and multiple hosts.

FIG. 3 illustrates an exemplary load balancing unit having separated functionality and an exemplary host.

FIG. 4 illustrates exemplary network load balancing infrastructure having separated classifying and forwarding functionality.

FIG. 5 is a flow diagram that illustrates an exemplary method for scaling out network load balancing infrastructure into different configurations.

FIG. 6 illustrates a first exemplary network load balancing infrastructure configuration from a device perspective.

FIG. 7 illustrates a second exemplary network load balancing infrastructure configuration from a device perspective.

FIGS. 8A and 8B illustrate first and second exemplary network load balancing infrastructure configurations from a component perspective.

FIGS. 9A and 9B illustrate first and second exemplary network load balancing infrastructure configurations from a resource perspective.

FIG. 10 illustrates an exemplary network load balancing approach that involves host status information.

FIG. 11 is a flow diagram that illustrates an exemplary method for network load balancing that involves host status information.

1 FIG. 12 illustrates an exemplary network load balancing approach that
2 involves health and load information.

3 FIG. 13A is an exemplary health and load table as illustrated in FIG. 12.

4 FIG. 13B is an exemplary consolidated health and load cache as illustrated
5 in FIG. 12.

6 FIG. 14 is a flow diagram that illustrates an exemplary method for network
7 load balancing that involves health and load information.

8 FIG. 15 illustrates an exemplary message protocol for communications
9 between the hosts and load balancing units that are illustrated in FIG. 12.

10 FIG. 16 illustrates an exemplary message transmission scheme for
11 communications between the hosts and load balancing units that are illustrated in
12 FIG. 12.

13 FIGS. 17A and 17B illustrate exemplary health and load information proxy
14 storage scenarios for health and load tables of FIG. 13A and for consolidated
15 health and load caches of FIG. 13B, respectively.

16 FIG. 18 illustrates an exemplary target host allotment procedure that utilizes
17 health and load information.

18 FIG. 19 illustrates an exemplary network load balancing approach that
19 involves session information.

20 FIG. 20 illustrates an exemplary network load balancing approach that
21 involves communicating session information using notifications and messages.

22 FIG. 21 is a flow diagram that illustrates an exemplary method for network
23 load balancing that involves communicating session information using
24 notifications and messages.

1 FIG. 22 illustrates an exemplary approach to managing session information
2 at multiple load balancing units.

3 FIG. 23A is an exemplary session table as illustrated in FIG. 20.

4 FIG. 23B is an exemplary distributed atom manager (DAM) table (DAMT)
5 as illustrated in FIG. 22.

6 FIG. 24 is a flow diagram that illustrates an exemplary method for
7 managing session information at multiple load balancing units.

8 FIG. 25 illustrates exemplary network load balancing infrastructure having
9 request routing functionality.

10 FIG. 26 is a flow diagram that illustrates an exemplary method for routing
11 incoming packets with regard to (i) session information and (ii) health and load
12 information.

13 FIG. 27 illustrates an exemplary traffic routing flow in the absence of
14 failures.

15 FIG. 28 illustrates an exemplary traffic routing flow in the presence of
16 failure(s).

17 FIG. 29 illustrates additional exemplary failover procedures for high
18 availability of network load balancing infrastructure.

19 FIG. 30 illustrates an exemplary operational implementation of traffic
20 routing interaction with health and load information.

21 FIG. 31 illustrates exemplary high availability mechanisms for network
22 load balancing infrastructure.

23 FIG. 32 illustrates an exemplary approach to application-level network load
24 balancing with connection migration.

25

1 FIG. 33 is a flow diagram that illustrates an exemplary method for
2 migrating a connection from a first device to a second device.

3 FIG. 34 illustrates an exemplary approach to connection migration from the
4 perspective of an originating device.

5 FIG. 35 illustrates an exemplary approach to connection migration from the
6 perspective of a targeted device.

7 FIG. 36 illustrates an exemplary approach to an offloading procedure for a
8 connection migration.

9 FIG. 37 illustrates an exemplary approach to an uploading procedure for a
10 connection migration.

11 FIG. 38 illustrates an exemplary approach to packet tunneling between a
12 forwarder and a host.

13 FIG. 39 is a flow diagram that illustrates an exemplary method for packet
14 tunneling between a first device and a second device.

15 FIG. 40 illustrates an exemplary computing (or general device) operating
16 environment that is capable of (wholly or partially) implementing at least one
17 aspect of network load balancing as described herein.

18

19 **DETAILED DESCRIPTION**

20 **Exemplary Network Load Balancing Paradigms**

21 This section describes exemplary paradigms for network load balancing
22 and is used to provide foundations, environments, contexts, etc. for the
23 descriptions in the following sections. This section primarily references FIGS. 1-
24 3.
25

1 FIG. 1 is an exemplary network load balancing paradigm 100 that illustrates
2 a load balancing infrastructure 106 and multiple hosts 108. Exemplary network
3 load balancing paradigm 100 includes multiple clients 102(1), 102(2) ... 102(m)
4 and multiple hosts 108(1), 108(2) ... 108(n), as well as network 104 and load
5 balancing infrastructure 106.

6 Each of clients 102 may be any device that is capable of network
7 communication, such as a computer, a mobile station, an entertainment appliance,
8 another network, and so forth. Clients 102 may also relate to a person and/or
9 entity that is operating a client device. In other words, clients 102 may comprise
10 logical clients that are users and/or machines. Network 104 may be formed from
11 one or more networks, such as the Internet, an intranet, a wired or wireless
12 telephone network, and so forth. Additional examples of devices for clients 102
13 and network types/topologies for network 104 are described below with reference
14 to FIG. 40 in the section entitled “Exemplary Operating Environment for
15 Computer or Other Device”.

16 Individual clients 102 are capable of communicating with one or more
17 hosts 108, and vice versa, across network 104 via load balancing infrastructure
18 106. Hosts 108 host one or more applications for interaction/communication with
19 clients 102, for use by clients 102, and so forth. Each host 108 may correspond to
20 a server and/or a device, multiple servers and/or multiple devices, part of a server
21 and/or part of a device, some combination thereof, and so forth. Particular
22 implementations for hosts 108 are described further below in the context of
23 different network load balancing situations. (However, back-end support for hosts
24 108 is generally not shown for the sake of clarity.) Furthermore, additional
25 examples of devices for hosts 108 are also described below with reference to FIG.

1 40 in the section entitled “Exemplary Operating Environment for Computer or
2 Other Device”.

3 Load balancing infrastructure 106 is reachable or locatable through network
4 104 at one or more virtual internet protocol (IP) addresses. Communications from
5 clients 102 (or other nodes) that are directed to the virtual IP address of load
6 balancing infrastructure 106 are received there and forwarded to a host 108. Load
7 balancing infrastructure 106 is comprised of hardware and/or software
8 components (not explicitly shown in FIG. 1).

9 Although load balancing infrastructure 106 is shown as an integral ellipse,
10 the infrastructure to effectuate load balancing may also be distributed to other
11 aspects of exemplary network load balancing paradigm 100. For example,
12 software component(s) of load balancing infrastructure 106 may be located at one
13 or more of hosts 108 as is described further below. Examples of architectures for
14 load balancing infrastructure 106 are described below with reference to FIG. 40 in
15 the section entitled “Exemplary Operating Environment for Computer or Other
16 Device”.

17 As indicated at (1), one or more of hosts 108 may provide host status
18 information from hosts 108 to load balancing infrastructure 106. This host status
19 information may be application specific. Examples of such host status information
20 are described further below and include health and/or load information, session
21 information, etc. for hosts 108. A particular implementation that includes
22 providing health and/or load information from hosts 108 to load balancing
23 infrastructure 106 is described below in the section entitled “Exemplary Health
24 and Load Handling”.

1 At (2), a request is sent from client 102(1) across network 104 to load
2 balancing infrastructure 106 at the virtual IP address thereof. The content, format,
3 etc. of a request from a client 102 may depend on the application to which the
4 request is directed, and the term “request” may implicitly include a response or
5 responses from host(s) 108, depending on the context. Kinds of client requests
6 include, but are not limited to:

7 1. Hyper text transfer protocol (HTTP) GET requests
8 from a client using a browser program. Depending on the
9 application (and more specifically, on the uniform resource locator
10 (URL) of the requests), it may be better to service the requests by
11 different sets of hosts, and the existence of a client “session” state on
12 the hosts may militate that requests from specific clients be routed to
13 specific hosts. The requests may be over a secure sockets layer
(SSL) (or other encrypted) connection.

14 2. Virtual private network (VPN) connections (e.g., the
15 hosts are a set of VPN servers). In this case, the “request” can be
16 considered to be a layer-2 tunneling protocol (L2TP) or point-to-
17 point tunneling protocol (PPTP) “connection” (the latter is a
18 combination of a transmission control protocol (TCP) control
19 connection and associated generic routing encapsulation (GRE) data
traffic).

20 3. Terminal server connections (e.g., the hosts are a set of
21 terminal servers).

22 4. Proprietary requests in the form of individual TCP
23 connections (one per request) employing a proprietary application-
specific protocol.

24 5. Simple object access protocol (SOAP) requests.
25

1 6. Real-time communication requests involving control
2 information over a TCP connection and latency-sensitive media
3 streaming over real-time protocol (RTP).

4 Thus, requests can take many diverse, application-specific forms. In certain
5 described implementations, load balancing infrastructure 106 may make
6 application-specific forwarding decisions.

7 At (3), load balancing infrastructure 106 forwards the request from 102(1)
8 to host 108(2) (in this example). Load balancing infrastructure 106 may consider
9 one or more of many factors when selecting a host 108 to which the request is to
10 be forwarded, depending on which implementation(s) described herein are being
11 employed. For example, load balancing infrastructure 106 may take into account:
12 the application health and/or load information of each host 108, session
13 information relating to client 102(1) as stored at a host 108, and so forth.

14 FIG. 2 is an exemplary network load balancing paradigm 200 that illustrates
15 multiple load balancing units 106 and multiple hosts 108. Specifically, load
16 balancing infrastructure 106 is shown as multiple load balancing units 106(1),
17 106(2) ... 106(u) in exemplary network load balancing paradigm 200.
18 Additionally, two router and/or switches 202(1) and 202(2) are illustrated.

19 Router/switches 202, if present, may be considered as part of or separate
20 from load balancing infrastructure 106 (of FIG. 1). Router/switches 202 are
21 responsible for directing overall requests and individual packets that are received
22 from network 104 to the shared virtual IP (VIP) address(es) of load balancing units
23 106. If a first router/switch 202 fails, the second router/switch 202 may takeover
24 for the first. Although two router/switches 202 are illustrated, one or more than
25 two router/switches 202 may alternatively be employed.

1 Router/switches 202 may be ignorant of the load balancing infrastructure or
2 load-balancing aware. If router/switches 202 are not load-balancing aware, one of
3 two exemplary options may be employed: For a first option, one load balancing
4 unit 106 is “assigned” the shared VIP address, and all network traffic is forwarded
5 thereto. This one load balancing unit 106 then evenly redistributes the traffic
6 across the other load balancing units 106. However, there are bottleneck and
7 failover issues with this first option (which can be mitigated if multiple VIP
8 addresses are shared and are split between multiple load balancing units 106). For
9 a second option, router/switches 202 are “tricked” into directing network traffic to
10 all load balancing units 106, which individually decide what traffic each should
11 accept for load balancing. However, there are inefficient effort duplication and
12 switch performance/compatibility issues with this second option.

13 If, on the other hand, router/switches 202 are load-balancing aware,
14 router/switches 202 can be made to distribute incoming network traffic
15 between/among multiple load balancing units 106 (e.g., in a round-robin fashion).
16 It should be understood that such load-balancing-aware routers/switches 202 are
17 capable of performing load balancing functions at a rudimentary level (e.g., in
18 hardware). For example, load-balancing-aware routers/switches 202 can perform
19 simple IP-address-based session affinity so that all packets from a specific source
20 IP address are directed to a same load balancing unit 106.

21 Each separately-illustrated load balancing unit 106 of load balancing units
22 106 may represent one physical device, multiple physical devices, or part of a
23 single physical device. For example, load balancing unit 106(1) may correspond
24 to one server, two servers, or more. Alternatively, load balancing unit 106(1) and
25 load balancing unit 106(2) may together correspond to a single server. An

1 exemplary load balancing unit 106 is described further below from a functional
2 perspective with reference to FIG. 3.

3 Two exemplary request paths [1] and [2] are illustrated in FIG. 2. For
4 request path [1], client 102(2) transmits a request over network 104 that reaches
5 router/switch 202(1). Router/switch 202(1) directs the packet(s) of the request that
6 originated from client 102(2) to load balancing unit 106(1). Load balancing unit
7 106(1) then forwards the packet(s) of the request to host 108(1) in accordance with
8 some load-balancing functionality (e.g., policy). For request path [2], client
9 102(m) transmits a request over network 104 that reaches router/switch 202(2).
10 Router/switch 202(2) directs the packet(s) of the request that originated from
11 client 102(m) to load balancing unit 106(u). Load balancing unit 106(u) then
12 forwards the packet(s) of the request to host 108(n) in accordance with some load-
13 balancing functionality. Exemplary load-balancing functionality is described
14 further below with reference to FIG. 3.

15 FIG. 3 illustrates an exemplary load balancing unit 106 having separated
16 functionality and an exemplary host 108. Load balancing unit 106 includes seven
17 (7) functional blocks 302-314. These functional blocks of load balancing unit 106
18 may be realized at least partially using software. Host 108 includes one or more
19 applications 316. In a described implementation, load balancing unit 106 includes
20 a forwarder 302, a classifier 304, a request router 306, a session tracker 308, a
21 connection migrator 310, a tunneler 312, and a health and load handler 314.

22 Health and load handler 314 is located partly at hosts 108 and partly on
23 devices of load balancing units 106. Health and load handler 314 monitors the
24 health and/or load (or more generally the status) of hosts 108 so that health and/or
25 load information thereof may be used for the load-balancing functionality (e.g.,

1 when making load-balancing decisions). Exemplary implementations for health
2 and load handler 314 are described further below, particularly in the section
3 entitled “Exemplary Health and Load Handling”.

4 Session tracker 308 may also be located partly at hosts 108 and partly on
5 devices of load balancing units 106. Session tracker 308 monitors sessions that
6 are established by clients 102 so that reconnections/continuations of previously-
7 established sessions may be facilitated by the load-balancing functionality. For
8 example, some applications keep application-specific client session data on the
9 hosts (which is also a type of host status information). These applications
10 typically expect that clients use the same host for the duration of any given
11 session. Exemplary types of sessions include: (i) a TCP connection (which is,
12 strictly speaking, a session); (ii) an SSL session; (iii) a secure IP (IPsec) session;
13 (iv) an HTTP cookie-based session; and so forth.

14 Although session tracker 308 is illustrated as a discrete block in load
15 balancing unit 106, session tracking functionality of session tracker 308 may
16 actually be implemented at a global level. In other words, session affinity is
17 supported across multiple load balancing units 106. Session tracker 308 includes a
18 centralized database and/or a distributed database of session information in order
19 to preserve session affinity. Exemplary implementations for session tracker 308,
20 with an emphasis on a distributed database approach, are described further below,
21 particularly in the section entitled “Exemplary Session Tracking”.

22 Classifier 304 uses the data acquired and maintained by health and load
23 handler 314 and/or session tracker 308, possibly in conjunction with other factors,
24 to classify incoming requests. In other words, classifier 304 selects a target host
25 108 for each incoming request from a client 102. Forwarder 302 forwards client

1 requests (and/or the packets thereof) in accordance with the targeted host 108 as
2 selected by classifier 304. Forwarder 302 and classifier 304 may operate on a per-
3 packet basis. Exemplary implementations for forwarder 302 and classifier 304 are
4 described further below, particularly in the sections entitled “Exemplary Approach
5 to Flexible Network Load Balancing” and “Exemplary Classifying, Forwarding,
6 and Request Routing”.

7 Request router 306, as contrasted with per-packet implementations of
8 forwarder 302 and classifier 304, can act as a proxy for an application running on
9 a host 108. For example, request router 306 may terminate TCP connections,
10 parse (perhaps partially) each logical request from a client 102, and resubmit each
11 logical request to the targeted host 108. Consequently, each logical request from a
12 client 102 may be directed to a different host 108, depending on the decisions
13 made by request router 306. Furthermore, request router 306 may perform pre-
14 processing on a connection (e.g., SSL decryption), may choose to absorb certain
15 requests (e.g., because request router 306 maintains a cache of responses), may
16 arbitrarily modify requests before forwarding them to hosts 108, and so forth.
17 Exemplary implementations for request router 306 are also described further
18 below, particularly in the sections entitled “Exemplary Approach to Flexible
19 Network Load Balancing” and “Exemplary Classifying, Forwarding, and Request
20 Routing”.

21 Connection migrator 310 enables a connection to be initially terminated at
22 load balancing unit 106 and then migrated such that the connection is subsequently
23 terminated at host 108. This connection migration can facilitate application-level
24 load balancing. Connection migrator 310 is capable of migrating a connection
25 from load balancing unit 106 to a host 108 in such a manner that that the original

1 termination at load balancing unit 106 is transparent to a requesting client 102 and
2 to applications 316 of the newly-terminating host 108. Tunneler 312 may utilize
3 an encapsulation scheme for the tunneling of packets that does not introduce an
4 overhead to each tunneled packet.

5 The functionality of tunneler 312 may also be used in situations that do not
6 involve a connection migration. Furthermore, connection migrator 310 and/or
7 tunneler 312 may additionally be used in non-load-balancing implementations.
8 Exemplary implementations for connection migrator 310, as well as for tunneler
9 312, are described further below, particularly in the section entitled “Exemplary
10 Connection Migrating with Optional Tunneling and/or Application-Level Load
11 Balancing”.

12 Any given implementation of a load balancing unit 106 may include one or
13 more of the illustrated functions. Although illustrated separately, each of the
14 functions of blocks 302-314 may actually be interrelated with, overlapping with,
15 and/or inclusive of other functions. For example, health and/or load information
16 of health and load handler 314 may be used by classifier 304. Also, connection
17 migrator 310 and tunneler 312 work in conjunction with forwarder 302 and
18 classifier 304. Certain other exemplary overlapping and interactions are described
19 herein below.

20 In a described implementation, host 108 runs and provides access to one or
21 more applications 316. Generally, applications 316 include file delivery programs,
22 web site management/server programs, remote access programs, electronic mail
23 programs, database access programs, and so forth. Specifically, applications 316
24 may include, but are not limited to, web servers such as Internet Information
25 Server® (IIS) from Microsoft® Corporation, terminal servers such as Microsoft®

1 Terminal Server™, and firewall and proxy products such as Internet Security and
2 Acceleration Server™ (ISA). Although the specific application 316 examples in
3 the preceding sentence relate to Microsoft® products, network load balancing as
4 described herein is not limited to any particular vendor(s), application(s), or
5 operating system(s).

6 **Exemplary Approach to Flexible Network Load Balancing**

7 This section illuminates how the network load balancing implementations
8 described in this and other sections herein provide a flexible approach to network
9 load balancing. This section primarily references FIGS. 4-9B.

10 As noted above, network load balancing functionality may be scaled up by
11 replacing a first network load balancer with a second, bigger and more powerful
12 network load balancer. The hardware capabilities of the second network load
13 balancer replicate the entirety of the hardware capabilities of the first network load
14 balancer, except that a greater capacity is provided. This is an inflexible approach
15 that can be very inefficient, especially when only one network load balancing
16 feature is limiting performance and precipitating an upgrade of a network load
17 balancer.

18 FIG. 4 illustrates exemplary network load balancing infrastructure having
19 separated classifying and forwarding functionality. The separated classifying
20 functionality and forwarding functionality are represented by classifier 304 and
21 forwarder 302, respectively. Although classifying and forwarding functions are
22 described further below, especially in the section entitled “Exemplary Classifying,
23 Forwarding, and Request Routing”, an initial description is presented here as an
24 example of interaction between network load balancing infrastructure functionality
25 and hosts 108.

1 In a described implementation, forwarder 302 corresponds to, and is the
2 network endpoint for, the virtual IP (VIP) address (or addresses). Forwarder 302
3 is a relatively low-level component that makes simplified and/or elementary policy
4 decisions, if any, when routing packets to a further or final destination. Forwarder
5 302 consults a routing table to determine this destination. Classifier 304 populates
6 the routing table based on one or more factors (e.g., host status information),
7 which are described further in other sections herein.

8 Clients 102 and hosts 108 also correspond to indicated network addresses.
9 Specifically, client 102(1) corresponds to address C1, client 102(2) corresponds to
10 address C2 ... client 102(m) corresponds to address Cm. Also, host 108(1)
11 corresponds to address H1, host 108(2) corresponds to address H2 ... host 108(n)
12 corresponds to address Hn.

13 Five communication paths (1)-(5) are shown in FIG. 4. Communication
14 path (1) is between client 102(1) and forwarder 302, and communication path (5)
15 is between forwarder 302 and host 108(1). Communication paths (2)-(4) are
16 between forwarder 302 and classifier 304. For simplicity in this example, the
17 connection associated with communication paths (1)-(5) is an HTTP TCP
18 connection. Furthermore, load balancing in this example relates to routing
19 incoming connections to the least loaded host 108, at least without any explicit
20 consideration of application-level load balancing.

21 Communication paths (1)-(5) indicate how forwarder 302 and classifier 304
22 load-balance a single HTTP TCP connection from client 102(1). At (1), client
23 102(1) initiates the TCP connection by sending a TCP SYN packet addressed to
24 the VIP address. The routing infrastructure of network 104 routes this packet to

1 forwarder 302 via router/switch 202(1), which is the “closest” router/switch 202 to
2 forwarder 302.

3 At (2), forwarder 302 consults a routing table, which may be internal to
4 forwarder 302 or otherwise accessible therefrom, in order to look up this
5 connection. This connection may be identified in the routing table by the TCP/IP
6 4-tuple (i.e., source IP address, source TCP port, destination IP address, destination
7 TCP port). Because this is the first packet of the connection, there is no entry in
8 the routing table. Forwarder 302 therefore applies the “default route” action,
9 which is to send this packet to classifier 304.

10 At (3), classifier 304 consults its (e.g., consolidated) cache of host status
11 information for hosts 108(1), 108(2) ... 108(n). Classifier 304 concludes that host
12 108(1) is available and the least loaded host 108 at this instant for this example.
13 Classifier 304 also “plumbs” a route in the routing table consulted by forwarder
14 302 for this TCP connection. For example, classifier 304 adds a route entry or
15 instructs forwarder 302 to add a route entry to the routing table that maps the TCP
16 connection (e.g., identified by the TCP 4-tuple) to a specific destination host 108,
17 which is host 108(1) in this example. More particularly, the route entry specifies
18 the network address H1 of host 108(1).

19 At (4), classifier 304 sends the TCP SYN packet back to forwarder 302.
20 Alternatively, classifier 304 may forward this initial TCP SYN packet to host
21 108(1) without using forwarder 302. Other options available to classifier 304 are
22 described further below.

23 At (5), forwarder 302 can access a route entry for the connection
24 represented by the SYN packet, so it forwards the packet to host 108(1) at address
25 H1. Forwarder 302 also forwards all subsequent packets from client 102(1) for

1 this connection directly to host 108(1). In other words, forwarder 302 can avoid
2 further interaction with classifier 304 for this connection. One or a combination of
3 mechanisms, which are described further below, may be used to delete the route
4 entry when the connection ceases.

5 For communication path (5) in many protocol environments, forwarder 302
6 cannot simply send the packets from client 102(1) as-is to host 108(1) at network
7 address H1 because these packets are addressed to the VIP address, which is
8 hosted by forwarder 302 itself. Instead, forwarder 302 may employ one or more
9 of the following exemplary options:

10 1. Forwarder 302 performs Network Address Translation
11 (NAT) by (i) overwriting the source (client 102(1)) IP address (C1)
12 and port number with the IP address and NAT-generated port number
13 of forwarder 302 and (ii) overwriting the destination IP address
14 (VIP) with the IP address (H1) of the host (108(1)).

15 2. Forwarder 302 performs “Half-NAT” by overwriting
16 the destination IP address (VIP) with the IP address (H1) of the host
17 (108(1)) so that the source (client 102(1)) IP address (C1) and port
18 number are preserved.

19 3. Forwarder 302 “tunnels” the packets received from
20 client 102(1) from forwarder 302 to host 108(1). Specifically in this
21 example, tunneling can be effectuated by encapsulating each packet
22 within a new IP packet that is addressed to host 108(1). Network-
23 load-balancing-aware software on host 108(1) reconstructs the
24 original packet as received at forwarder 302 from client 102(1).
25 This original packet is then indicated up on a virtual interface at host

1 108(1) (e.g., the VIP address corresponding to forwarder 302 is
2 bound to this virtual interface at host 108(1)). Exemplary
3 implementations of such tunneling are described further below with
4 reference to tunneler 312, especially for connection migration
5 scenarios and particularly in the section entitled “Exemplary
6 Connection Migrating with Optional Tunneling and/or Application-
7 Level Load Balancing”.

8 Although FIGS. 4-9B show two specific separated functions, namely
9 classifying and forwarding, it should be understood that other functions, such as
10 those of request router 306, session tracker 308, connection migrator 310, and
11 health and load handler 314, may also be scaled out independently (e.g., factored
12 out independently), as is described further below. Furthermore, it should be noted
13 that one or more than two functions may be separated and scaled out
14 independently at different times and/or simultaneously. Also, although TCP/IP is
15 used for the sake of clarity in many examples in this and other sections, the
16 network load balancing principles described herein are applicable to other
17 transmission and/or communication protocols.

18 In the exemplary manner of FIG. 4, network load balancing functions (such
19 as those shown in FIG. 3) may be separated from each other for scalability
20 purposes. They may also be separated and duplicated into various configurations
21 for increased availability. Exemplary configurations for scalability and/or
22 availability are described below with reference to FIGS. 6-9B after the method of
23 FIG. 5 is described.

24 FIG. 5 is a flow diagram 500 that illustrates an exemplary method for
25 scaling out network load balancing infrastructure into different configurations.

1 Flow diagram 500 includes three blocks 502-506. Although the actions of flow
2 diagram 500 may be performed in other environments and with a variety of
3 software schemes, FIGS. 1-4 and 6-9B are used in particular to illustrate certain
4 aspects and examples of the method.

5 At block 502, network load balancing infrastructure is operated in a first
6 configuration. For example, each configuration may relate to one or more of a
7 selection, proportion, and/or interrelationship of different load balancing
8 functionalities; a number of and/or type(s) of different devices; an organization
9 and/or layout of different components; a distribution and/or allocation of
10 resources; and so forth. At block 504, the network load balancing infrastructure is
11 scaled out. For example, separated load balancing functionalities may be
12 expanded and/or concomitantly contracted on an individual and/or independent
13 basis. At block 506, the scaled out network load balancing infrastructure is
14 operated in a second configuration.

15 As noted above, a monolithic network load balancer may be scaled up by
16 increasing network load balancing functionality in its entirety by supplanting
17 previous network load balancing hardware with more-powerful network load
18 balancing hardware. In contradistinction, scaling out network load balancing
19 infrastructure can enable network load balancing (sub-)functions to be scaled out
20 individually and/or independently. It can also enable network load balancing
21 functions to be scaled out together or individually between and among different
22 numbers of devices. Device, component, and resource-oriented scaling out
23 examples are provided below.

24 FIG. 6 illustrates a first exemplary network load balancing infrastructure
25 configuration from a device perspective. In this first device-oriented network load

1 balancing infrastructure configuration, three devices 602(1), 602(2), and 602(3)
2 are illustrated. However, one, two, or more than three devices 602 may
3 alternatively be employed.

4 As illustrated, a forwarder 302(1), a classifier 304(1), and a host 108(1) are
5 resident at and executing on device 602(1). A forwarder 302(2), a classifier
6 304(2), and a host 108(2) are resident at and executing on device 602(2). Also, a
7 forwarder 302(3), a classifier 304(3), and a host 108(3) are resident at and
8 executing on device 602(3). Thus, in this first device-oriented network load
9 balancing infrastructure configuration, a respective forwarder 302, classifier 304,
10 and host 108 are sharing the resources of each respective device 602.

11 In operation, forwarders 302 are the network endpoints for the VIP
12 address(es). Any classifier 304 may plumb a route for a connection to any host
13 108, depending on host status information. For example, classifier 304(2) may
14 plumb a route for a new incoming connection to host 108(3). In accordance with a
15 new route entry for this connection, forwarder 302(2) forwards subsequent packets
16 to host 108(3).

17 In one alternative device-oriented network load balancing infrastructure
18 configuration to which the illustrated first one may be scaled out, a fourth device
19 602(4) (not explicitly shown in FIG. 6) may be added that includes a forwarder
20 302(4), a classifier 304(4), and a host 108(4). If, on the other hand, sufficient
21 classification functionality is already present with classifiers 304(1-3) but
22 additional forwarding functionality can benefit the request handling of hosts 108, a
23 fourth device 602(4) may be added that includes a forwarder 302(4) and optionally
24 a host 108(4). For this scaled-out configuration, another classifier 304(1, 2, or 3)

1 may plumb routes for forwarder 302(4) to any of hosts 108(1, 2, or 3) and host
2 108(4), if present.

3 The first device-oriented exemplary network load balancing infrastructure
4 configuration of FIG. 6 may be especially appropriate for smaller hosting
5 situations in which separate devices for the network load balancing infrastructure
6 are not technically and/or economically worthwhile or viable. However, as the
7 hosting duties expand to a greater number (and/or a greater demand on the same
8 number) of hosts 108 or if the network load on hosts 108 is significant, the first
9 device-oriented exemplary network load balancing infrastructure configuration
10 may be scaled out to accommodate this expansion, as represented by a second
11 device-oriented exemplary network load balancing infrastructure configuration of
12 FIG. 7.

13 FIG. 7 illustrates a second exemplary network load balancing infrastructure
14 configuration from a device perspective. In this second device-oriented network
15 load balancing infrastructure configuration, three devices 602(1), 602(2), and
16 602(3) are also illustrated. Again, one, two, or more than three devices 602 may
17 alternatively be employed.

18 As illustrated, forwarder 302(1) and classifier 304(1) are resident at and
19 executing on device 602(1). Forwarder 302(2) and classifier 304(2) are resident at
20 and executing on device 602(2). Also, forwarder 302(3) and classifier 304(3) are
21 resident at and executing on device 602(3). Thus, in this second device-oriented
22 network load balancing infrastructure configuration, each respective forwarder
23 302 and classifier 304 are not sharing the resources of each respective device 602
24 with a host 108. Furthermore, the network load balancing infrastructure may be
25 servicing any number of hosts 108.

1 In operation, forwarders 302 are again the network endpoints for the VIP
2 address(es). Also, any classifier 304 may plumb a route for a connection to any
3 host 108, depending on host status information. For example, classifier 304(3)
4 may plumb a route for a new incoming connection to host 108(2). In accordance
5 with a new route entry for this connection, forwarder 302(3) forwards subsequent
6 packets to host 108(2).

7 Hence, network load balancing infrastructure as realized in software, for
8 example, may be scaled out by moving the network load balancing infrastructure
9 (or part thereof) from devices that are shared with hosts 108 to devices that are not
10 shared with hosts 108. Also, as alluded to above for FIG. 6, another device 602(4)
11 may be added to the network load balancing infrastructure to provide additional
12 forwarding functionality, additional classifying functionality, additional
13 functionality of both types, and so forth.

14 FIGS. 8A and 8B illustrate first and second exemplary network load
15 balancing infrastructure configurations from a component perspective. As
16 illustrated, first component-oriented exemplary network load balancing
17 infrastructure configuration 800 includes four components. Second component-
18 oriented exemplary network load balancing infrastructure configuration 850
19 includes six components. An alternative second configuration 850 includes a
20 seventh component as indicated by the dashed-line block, which is described
21 further below.

22 Specifically, first component-oriented exemplary network load balancing
23 infrastructure configuration 800 (or first configuration 800) includes (i) two
24 forwarders 302(1) and 302(2) and (ii) two classifiers 304(1) and 304(2). Second
25 exemplary component-oriented network load balancing infrastructure

1 configuration 850 (or second configuration 850) includes (i) four forwarders
2 302(1), 302(2), 302(3), and 302(4) and (ii) two classifiers 304(1) and 304(2).
3 Thus, first configuration 800 is scaled out to second configuration 850 by adding
4 two components, which are forwarding components in this example.

5 In a described implementation, each respective network-load-balancing-
6 related functional component corresponds to a respective device (not explicitly
7 shown in FIG. 8A or 8B); however, each component may alternatively correspond
8 to part of a device or more than one device. For example, forwarders 302(1) and
9 302(2) may be distributed across three devices. Or forwarder 302(1) and classifier
10 304(1) may correspond to a first device, and forwarder 302(2) and classifier
11 304(2) may correspond to a second device.

12 Two network-load-balancing-related functional components are added to
13 scale out first configuration 800 to second configuration 850. However, one
14 component (or more than two) may alternatively be added to scale out the network
15 load balancing infrastructure. Furthermore, two or more different types of
16 functional components may be scaled out “simultaneously”. For example, as
17 illustrated by the dashed-line block, another classifying component (e.g., classifier
18 304(3)) may also be added when scaling out first configuration 800 to second
19 configuration 850.

20 Moreover, scaling by two or more different types of functional components
21 may be performed in similar (e.g., equivalent) or dissimilar proportions to each
22 other. As illustrated, adding forwarder components 302(3) and 302(4) while not
23 adding any classifier component 304 or while adding a single classifier component
24 304(3) represent a scaling out at dissimilar proportions. However, two classifier
25 components 304(3) and 304(4) (the latter of which is not explicitly illustrated in

1 FIG. 8B) may be added while the two forwarder components 302(3) and 302(4)
2 are added for a scaling out at similar proportions. Regardless, each individual
3 network-load-balancing-related functional component may consume a different
4 amount of the available network load balancing infrastructure resources, as is
5 described with reference to FIGS. 9A and 9B.

6 FIGS. 9A and 9B illustrate first and second exemplary network load
7 balancing infrastructure configurations from a resource perspective. First
8 resource-oriented exemplary network load balancing infrastructure configuration
9 900 (or first configuration 900) includes a first resource distribution or allocation
10 for a load balancing unit 106. Second resource-oriented exemplary network load
11 balancing infrastructure configuration 950 (or second configuration 950) includes
12 a second resource distribution for load balancing unit 106.

13 As illustrated, first configuration 900 includes a 70%-30% resource
14 distribution, and second configuration 950 includes a 40%-60% resource
15 distribution. Such resources may include total device resources (e.g., number of
16 devices), processing resources (e.g., number of processor cycles), memory
17 resources (e.g., portion of cache, main memory, etc.), network bandwidth and/or
18 interface resources (e.g., bits per second and/or physical network interface cards
19 (NICs)), and so forth.

20 Specifically for first configuration 900, forwarder 302 consumes 70% of
21 the resources of load balancing unit 106 while classifier 304 consumes 30% of
22 these resources. After reallocation during a scaling out procedure to produce
23 second configuration 950, forwarder 302 consumes 40% of the resources of load
24 balancing unit 106 while classifier 304 consumes 60% of these resources.
25

1 In an exemplary situation, first configuration 900 might facilitate better
2 network load balancing performance when fewer, longer transactions are being
3 handled by the associated hosts (not shown in FIGS. 9A and 9B) because
4 classification functionality is utilized upon initial communication for a connection
5 and forwarding functionality is utilized thereafter. Second configuration 950, on
6 the other hand, might facilitate better network load balancing performance when
7 more, shorter transactions are being handled by the associated hosts because the
8 classification functionality is utilized for a greater percentage of the total number
9 of packets funneled through the network load balancing infrastructure. In this
10 situation, if request routing functionality is also being employed, then request
11 router(s) 306 are also allocated a percentage of the total computing resources. The
12 resource distribution among the three functionalities may be adjusted while
13 handling connections (e.g., adjusted “on the fly”) depending on current resource
14 consumption and/or deficits.

15 As indicated above with reference to FIGS. 2 and 3, each load balancing
16 unit 106 may correspond to all or a part of a total network load balancing
17 infrastructure 106. For any given physically, logically, arbitrarily, etc. defined or
18 stipulated load balancing unit 106, the resources thereof may be re-allocated
19 during a scale out procedure. More specifically, a resource distribution
20 between/among different network-load-balancing-related separated functions of a
21 load balancing unit 106 may be altered in a scale out procedure. Furthermore,
22 more than two different functions, as well as other network-load-balancing-related
23 functions that are not specifically illustrated in FIGS. 9A and 9B, may be allocated
24 differing resource percentages.

1 The percentage of total system resources allocated to all load balancing
2 functions may also be altered in a scale out procedure. As a general processing
3 power example, the percentage of total processing power that is devoted to load
4 balancing may be gradually increased as the amount of traffic that needs to be load
5 balanced increases.

6 Network load balancing software may optionally perform monitoring to
7 analyze and determine whether resources should be reallocated. For example, the
8 network load balancing software may monitor the processor utilization of different
9 network-load-balancing-related functions. The actual reallocation may also
10 optionally be performed automatically by the network load balancing software in
11 an offline or online mode.

12 It should be understood that a scaling out capability of network load
13 balancing infrastructure (e.g., as realized at least partially in software) as described
14 herein may relate to different installations and not necessarily a change to a single
15 installation. In a resource-oriented example, network load balancing infrastructure
16 as described herein may be configured in accordance with one resource
17 distribution in one installation environment and may be configured in accordance
18 with another different resource distribution in another installation environment
19 having different operational parameters. Additionally, the capabilities, features,
20 options, etc. described above with regard to scaling out are also applicable for
21 “scaling in”. In other words, resources devoted to network load balancing
22 infrastructure (or sub-functions thereof) may also be reduced.

23 **Exemplary Health and Load Handling**

24 This section describes how host status information, such as health and/or
25 load information, may be collected for and utilized in network load balancing.

1 This section primarily references FIGS. 10-18 and illuminates health and load
2 functionality such as that provided by health and load handler 314 (of FIG. 3). As
3 described above with reference to FIG. 3, each host 108 hosts one or more
4 applications 316. Health and load handler 314 utilizes health and/or load
5 information that relates to applications 316 and/or hosts 108 for certain described
6 implementations of network load balancing.

7 FIG. 10 illustrates an exemplary network load balancing approach that
8 involves host status information (HSI) 1006. Each host 108(1), 108(2) ... 108(n)
9 includes one or more applications 316(1), 316(2) ... 316(n), respectively. These
10 hosts 108 generally and these applications 316 specifically may change statuses
11 from time to time.

12 For example, hosts 108 and applications 316 may be accepting new
13 connections or not accepting new connections. Also, they may be quickly
14 handling client requests or slowly handling client requests. Furthermore, they may
15 have many resources in reserve or few unused resources. All or any part of such
16 data, or other data, may comprise host status information 1006. Generally, host
17 status information 1006 provides an indication of the status of some aspect of
18 hosts 108 and/or applications 316 that are running thereon.

19 In a described implementation, each host 108(1), 108(2) ... 108(n) includes
20 a host status information (HSI) determiner 1002(1), 1002(2) ... and 1002(n),
21 respectively. Each host 108(1), 108(2) ... 108(n) also includes a host status
22 information (HSI) disseminator 1004(1), 1004(2) ... and 1004(n), respectively.
23 Each host status information determiner 1002 and/or host status information
24 disseminator 1004 may be part of load balancing infrastructure (LBI) 106.

25

1 Each host status information determiner 1002 determines host status
2 information 1006 for its respective host 108 and/or applications 316 that are
3 running thereon. Exemplary techniques for determining such host status
4 information 1006 are described below with reference to FIGS. 12-14, and
5 particularly FIG. 13A. Each host status information disseminator 1004
6 disseminates host status information 1006 for its respective host 108 and/or
7 applications 316 to load balancing infrastructure 106 (e.g., those portion(s) of load
8 balancing infrastructure 106 that are not located on hosts 108). Exemplary
9 techniques for disseminating such host status information 1006 are described
10 below with reference to FIGS. 12-17, and particularly FIGS. 13B and 15-17.

11 Specifically, each host status information disseminator 1004 disseminates
12 host status information 1006 (directly or indirectly) to each load balancing unit
13 (LBU) 106 of load balancing infrastructure 106 that includes at least one health
14 and load handler 314 and/or classifier 304. Load balancing infrastructure 106
15 refers to host status information 1006 when implementing network load balancing.
16 For example, as indicated by logic 1008, load balancing infrastructure 106 is
17 capable of making load balancing decisions responsive to host status information
18 1006.

19 In operation at (1), host status information determiners 1002 determine host
20 status information 1006 for respective hosts 108 and/or applications 316. At (1)
21 and (2), host status information disseminators 1004 disseminate host status
22 information 1006 from hosts 108 to load balancing infrastructure 106. For
23 example, host status information 1006 may be disseminated to individual load
24 balancing units 106. At (3), logic 1008 makes network load balancing decisions

1 responsive to host status information 1006. At (4), connections are forwarded to
2 targeted hosts 108 based on these network load balancing decisions.

3 FIG. 11 is a flow diagram 1100 that illustrates an exemplary method for
4 network load balancing that involves host status information. Flow diagram 1100
5 includes three blocks 1102-1106. Although the actions of flow diagram 1100 may
6 be performed in other environments and with a variety of software schemes, FIGS.
7 1-3 and 10 are used in particular to illustrate certain aspects and examples of the
8 method.

9 At block 1102, host status information is sent from hosts to load balancing
10 units. For example, host status information 1006 may be sent from hosts 108 to
11 load balancing units 106. At block 1104, the host status information is received
12 from the hosts at the load balancing units. For example, load balancing units 106
13 may receive host status information 1006 from hosts 108. At block 1106, load
14 balancing decisions are made responsive to the received host status information.
15 For example, logic 1008 at load balancing units 106 may make decisions for
16 network load balancing responsive to host status information 1006.

17 Thus in FIG. 10, load balancing infrastructure 106 collects host status
18 information 1006 from hosts 108 (and/or applications 316 thereof) and load
19 balances incoming requests that are directed to hosts 108 responsive to host status
20 information 1006. As described further below with reference to FIGS. 12-18, this
21 host status information 1006 may be application-specific. As also described
22 further below, examples of host status information 1006 include health and/or load
23 information.

24 FIG. 12 illustrates an exemplary network load balancing approach that
25 involves health and/or load information (HLI) 1206. Hosts 108(1), 108(2) ...

1 108(n) are coupled to load balancing units 106(1), 106(2) ... 106(u) via a
2 communication linkage 1210 such as a network.

3 As illustrated, hosts 108 communicate health and load information 1206 to
4 load balancing units 106 using communication linkage 1210. The bi-directional
5 communication of health and load information 1206, as indicated by the double-
6 pointed arrow, refers to a two-way communication from load balancing units 106
7 to hosts 108 that provides certain completeness, coherency, correctness, etc. such
8 that hosts 108 and/or load balancing units 106 may fail independently of one
9 another. Such two-way communications from load balancing units 106 to hosts
10 108 are described further below with particular reference to FIG. 15.

11 Health information reflects whether a given host and/or application is
12 capable of handling client requests. Load information reflects the number,
13 amount, and/or level of client requests that the given host and/or application is
14 capable of handling at a particular moment. In other words, load can reflect
15 directly and/or inversely an available number, amount, and/or level of total
16 capacity of the given host and/or application. As noted above, implementations
17 described with reference to FIGS. 12-18 focus on health and/or load information;
18 however, those implementations are also applicable to general status information
19 for hosts (including the applications thereof).

20 In a described implementation, each host 108(1), 108(2) ... 108(n) includes
21 a respective health and load infrastructure (H&LI) component 1202(1), 1202(2) ...
22 1202(n). Each health and load infrastructure component 1202 may optionally be a
23 portion of load balancing infrastructure 106 that is resident at and executing on
24 each host 108. Health and load information 1206 may be realized in software.
25 When functioning, each health and load infrastructure 1202(1), 1202(2) ...

1 1202(n) creates and maintains a respective health and load (H&L) table 1204(1),
2 1204(2) ... 1204(n).

3 These health and load tables 1204 may include application-specific entries.
4 Health and load information 1206 that is stored in health and load tables 1204 may
5 be independent of load balancing infrastructure 106. For example, administrators,
6 designers, etc. may specify criteria for health and load information 1206 at
7 configuration time. Additionally, entities external to a device that is or that has a
8 host 108 may contribute to determining health and load information 1206 for
9 applications 316 on the device. An exemplary health and load table 1204 is
10 described further below with reference to FIG. 13A.

11 Each load balancing unit 106(1), 106(2) ... 106(u) includes a respective
12 consolidated health and load (H&L) cache 1208(1), 1208(2) ... 1208(u). Each
13 consolidated health and load cache 1208 includes the information from each health
14 and load table 1204(1), 1204(2) ... 1204(n). Consequently, each load balancing
15 unit 106 is provided with quick (e.g., cached) access to health and load
16 information 1206 for each host 108 for which load balancing units 106 are load
17 balancing network traffic.

18 In operation, health and load infrastructures 1202 push health and load
19 information 1206 from health and load tables 1204 to consolidated health and load
20 caches 1208. The mechanism to provide health and load information 1206 is
21 event driven such that changes to health and load tables 1204 are provided to
22 consolidated health and load caches 1208 in a timely, scaleable manner.

23 FIG. 13A is an exemplary health and load table 1204 as illustrated in FIG.
24 12. In a described implementation, health and load table 1204 includes multiple
25 entries 1302 that are each associated with a different application 316. Each entry

1 1302 may correspond to a row in health and load table 1204 that has three
2 columns. These columns correspond to application identifier (ID) 1302(A),
3 application status characterization 1302(B), and load balancer directive 1302(C).

4 Because each entry 1302 is associated with a particular application 316, a
5 row is added as each application is spun up (e.g., by an administrator). Likewise,
6 a row is deleted/removed each time an application is closed down. Similarly,
7 individual fields in columns 1302(A), 1302(B), and/or 1302(C) are
8 modified/updated when a value thereof changes. For example, when a status
9 characterization value changes for a given application 316, a value in a field of
10 application status characterization 1302(B) for entry 1302 of the given application
11 316 is updated.

12 The additions and deletions of entries 1302 for applications 316 may be
13 effectuated with input from a control manager at the host 108. For example, a
14 control manager portion of an operating system knows when an application 316 is
15 started and stopped because it is actively involved in the starting and stopping of
16 applications 316. Hence, a control manager may identify that it has, at least in
17 part, started an application 316, and the control manager may establish that it has,
18 at least in part, stopped the application 316. Health and load infrastructure 1202
19 may therefore be informed of the starting and stopping of applications 316 by the
20 control manager. Hence, no such explicit communication from applications 316
21 has to be provided to health and load infrastructure 1202. An example of a control
22 manager is the Service Control Manager (SCM) of the Windows® Operating
23 System from Microsoft® Corporation.

24 Application identifier 1302(A) includes information that is used to uniquely
25 identify the application 316 to which entry 1302 is associated. Application

1 identifier 1302(A) may include one or more of the following for the associated
2 application 316: the virtual IP address and port, the physical IP address and port,
3 the protocol used, and any protocol-specific information. The protocol may be
4 HTTP, IPsec, SOAP, and so forth. The protocol-specific information may be a
5 URL pattern or string to further delineate the application associated with entry
6 1302. Thus, application identifier 1302(A) more particularly refers to a specific
7 application endpoint on a particular host 108.

8 Other application identifiers may alternatively be employed. For example,
9 to reduce communication bandwidth, application identifier 1302(A) may be a 32-
10 bit number that maps to the above exemplary information at health and load
11 infrastructure 1202 and at load balancing units 106. Moreover, any of the fields in
12 entry 1302 may actually contain a globally unique identifier (GUID) that is used
13 as a key to look up the true information for the field.

14 Application status characterization 1302(B) includes information that
15 reflects the status of the application 316 to which entry 1302 is associated.
16 Application status characterization 1302(B) includes the following for the
17 associated application 316: application health, application load, and application
18 capacity. Application health is a quasi-Boolean value that indicates whether an
19 application is functioning. Application health can be healthy, failing, or unknown.
20 Application health is a relatively-instantaneous value and is communicated with
21 relatively low latency (e.g., of approximately a second or a few seconds) to load
22 balancing units 106 when the application health value changes.

23 Application load is a value that indicates how occupied or busy a given
24 application is and thus, directly or inversely, how much additional load the given
25 application can handle. Application load is a relatively slowly-changing or

1 averaged value that can be smoothed with a hysteresis-inducing mechanism, if
2 desired, to eliminate transient spikes of increased or decreased load. It is
3 communicated relatively infrequently to load balancing units 106 (e.g.,
4 approximately one to four times a minute). The value of application load is given
5 meaning with regard to application capacity.

6 Application capacity is a value that indicates the maximum capacity of the
7 application. It is selected in a generic manner to be meaningful for a given context
8 but still sufficiently flexible for other contexts. Application capacity is a unit-less,
9 bounded number (e.g., 0-99) that is determinable at configuration time. It may be
10 based on processing power, memory size/speed, network access, some
11 combination thereof, and so forth. Application capacity expresses relative
12 capacities between and among other applications of the same type in a set of hosts
13 108(1, 2 ... n).

14 Thus, relative to application capacity, application load gains meaning.
15 Application load for a given application is a percentage of application capacity for
16 the given application. Alternatively, application load can be expressed as a unit-
17 less number from which the percentage can be ascertained in conjunction with the
18 value of application capacity.

19 Load balancer directive 1302(C) includes information that reflects the
20 desired and/or expected state of the directive established by health and load
21 infrastructure 1202 for load balancing units 106 with respect to an application 316
22 to which entry 1302 is associated. Load balancer directive 1302(C) includes the
23 following for the associated application 316: target load balancing state and
24 current load balancing state.

25

1 The target load balancing state reflects the state of the directive to load
2 balancing units 106 as desired by health and load infrastructure 1202. The current
3 load balancing state reflects what health and load infrastructure 1202 understands
4 the current state of the directive to load balancing units 106 to be as recorded at
5 load balancing units 106. The current load balancing state thus reflects the load
6 balancing directive that health and load infrastructure 1202 expects load balancing
7 units 106 to be currently operating under as dictated using a communication
8 protocol. Such an exemplary communication protocol is described further below
9 with reference to FIG. 15. The interaction and relationship between the target load
10 balancing state and the current load balancing state is also further clarified with
11 the description of FIG. 15.

12 The target load balancing state and the current load balancing state may
13 each take a value of active, inactive, or draining. An active directive indicates that
14 new requests/connections are welcome and may be targeted at the application that
15 is associated with entry 1302. An inactive directive indicates that no additional
16 packets should be forwarded to the associated application. A draining directive
17 indicates that no packets for new requests/connections should be sent to the
18 associated application but that packets for existing requests/connections should
19 continue to be forwarded to the associated application.

20 In a described implementation, the definitive version of respective health
21 and load information 1206 is stored at health and load tables 1204 that are located
22 at each respective host 108 of multiple hosts 108. With this implementation, if a
23 host 108 crashes, the health and load information 1206 that is lost pertains to those
24 applications 316 that also crashed. A measure of high availability is therefore
25 gained automatically without duplicating data. However, the definitive version of

1 health and load information 1206 may alternatively be stored elsewhere. Other
2 such storage options include load balancing units 106 themselves, a host 108 that
3 (as its sole task or along with hosting duties) stores and maintains health and load
4 information 1206 for multiple other (including all other) hosts 108, another
5 separate and/or external device, and so forth.

6 If the definitive version of health and load information 1206 is stored and
7 maintained elsewhere besides being distributed across hosts 108(1, 2 ...n), such
8 health and load information 1206 may be stored redundantly (e.g., also stored in a
9 duplicative device, backed-up, etc.) for high-availability purposes. Exemplary
10 proxy scenarios for storing health and load information 1206 are described below
11 with reference to FIGS. 17A and 17B. FIG. 17A is directed to a proxy scenario for
12 health and load tables 1204, and FIG. 17B is directed to a proxy scenario for
13 consolidated health and load caches 1208.

14 FIG. 13B is an exemplary consolidated health and load cache 1208 as
15 illustrated in FIG. 12. In a described implementation, each consolidated health and
16 load cache 1208 in each load balancing unit 106 includes at least part of the
17 information stored in each health and load table 1204 for each health and load
18 infrastructure 1202 at each host 108. The cached health and load information may
19 be organized in any manner in consolidated health and load cache 1208.

20 As illustrated, consolidated health and load cache 1208 includes a cache for
21 each host 108(1), 108(2) ... 108(n) that replicates part or all of the information in
22 the health and load table 1204 of each respective host 108(1, 2 ... n). Specifically,
23 consolidated health and load cache 1208 includes a cache for host #1 1304(1), a
24 cache for host #2 1304(2) ... a cache for host #n 1304(n). Thus, the illustrated
25 consolidated health and load cache 1208 is organized at a broad level by host

1 108(1, 2 ... n), with each individual cache 1304 including application-specific
2 entries for the corresponding respective host 108(1, 2 ... n). Alternatively,
3 consolidated health and load cache 1208 may be organized at a broad level by type
4 of application 316, with individual blocks that are directed to a specific application
5 type further divided by host 108(1, 2 ... n). Other data structure formats may also
6 be employed.

7 FIG. 14 is a flow diagram that illustrates an exemplary method for network
8 load balancing that involves health and load information. Flow diagram 1400
9 includes eight blocks 1402-1416. Although the actions of flow diagram 1400 may
10 be performed in other environments and with a variety of software schemes, FIGS.
11 1-3 and 12-13B are used in particular to illustrate certain aspects and examples of
12 the method. For example, the actions of two blocks 1402-1404 are performed by a
13 host 108, and the actions of six blocks 1406-1416 are performed by a load
14 balancing unit 106.

15 At block 1402, health and load information at a host is determined. For
16 example, health and load information 1206 for applications 316(2) may be
17 ascertained by health and load infrastructure 1202(2) and stored in health and load
18 table 1204(2) at host 108(2). At block 1404, the determined health and load
19 information is disseminated to load balancing units. For example, health and load
20 infrastructure 1202(2) may send health and load information 1206 for applications
21 316(2) to load balancing units 106(1, 2 ... u). As indicated by arrow 1418, the
22 actions of blocks 1402 and 1404 are repeated so that (application) health and load
23 may be continually monitored and updated as changes occur.

24 At block 1406, health and load information is received from hosts. For
25 example, load balancing unit 106(1) may receive health and load information 1206

1 from multiple hosts 108(1, 2 ... n), which includes health and load information
2 1206 for applications 316(2) of host 108(2). At block 1408, the received health
3 and load information is cached. For example, load balancing unit 106(1) may
4 store health and load information 1206 from hosts 108(1, 2 ... n) into consolidated
5 health and load cache 1208(1). With reference to the FIG. 13B implementation of
6 a consolidated health and load cache 1208(1), health and load information 1206
7 for applications 316(2) from host 108(2) may be stored in cache for host #2
8 1304(2). As indicated by arrow 1420, the actions of blocks 1406 and 1408 are
9 repeated so that (application) health and load information may be continually
10 received and updated as changes occur.

11 As indicated by dashed arrow 1422, load balancing units 106 are also
12 handling communications from clients 102 while handling (application) health and
13 load issues. At block 1410, a packet requesting a new connection is received. For
14 example, load balancing unit 106(1) may receive a TCP SYN packet from client
15 102(2) through network 104. At block 1412, the cached health and load
16 information is consulted. For example, load balancing unit 106(1) may consult
17 consolidated health and load cache 1208(1). More particularly, load balancing unit
18 106(1) may consult entries that are associated with the application to which the
19 TCP SYN packet is directed across caches for hosts #1, #2 ... #n 1304(1, 2 ... n).

20 At block 1414, a host is selected responsive to the cached health and load
21 information. For example, load balancing unit 106(1) may select host 108(2)
22 having application(s) 316(2) responsive to health and load information 1206 that is
23 cached in consolidated health and load cache 1208(1). The selected application
24 316 (and host 108) should be healthy and able to accept additional load (e.g.,
25

1 possibly the least loaded application among those applications that are of the
2 application type to which the TCP SYN packet is directed).

3 The consulting of the cached health and load information (at block 1412)
4 and the host-selecting responsive to the cached health and load information (at
5 block 1414) may be performed prior to reception of a specific new-connection-
6 requesting packet and/or using a batched scheme. Also, the selecting may be in
7 accordance with any of many schemes. For example, a token based or a round-
8 robin based scheme may be employed. With either scheme, the selection may
9 involve a weighting of relative loads among the application options. This
10 consultation and selection, along with the token and round-robin based schemes,
11 are described further below with reference to FIG. 18 and in the section entitled
12 “Exemplary Classifying, Forwarding, and Request Routing”, especially with
13 regard to classifying functionality.

14 After the target host is selected at block 1414, the new-connection-
15 requesting packet may be sent thereto. At block 1416, the packet received from
16 the client is forwarded to the selected host. For example, the TCP SYN packet is
17 forwarded from load balancing unit 106(1) to selected host 108(2). The
18 forwarding of this initial packet may be effectuated directly by a classifier 304 or
19 by a forwarder 302, as is also described further below in the section entitled
20 “Exemplary Classifying, Forwarding, and Request Routing”.

21 For a described implementation, health and load infrastructure 1202 is
22 resident at and distributed across multiple hosts 108 as well as being located at
23 load balancing units 106 (as represented by health and load handler 314). Health
24 and load infrastructure 1202 has three responsibilities. First, it exposes listening
25 point(s) to attain application status updates for application status characterizations

1 1302(B) of health and load tables 1204. Second, it synthesizes the application
2 status information to determine what load balancing units 106 should do, which is
3 embodied in load balancer directive 1302(C). Third, health and load infrastructure
4 1202 communicates this directive from hosts 108 to load balancing units 106.

5 The directive content of load balancer directive 1302(C) is effectively a
6 digested version of the information for application status characterizations
7 1302(B). However, load balancing units 106 may also receive the raw information
8 of application status characterizations 1302(B) as well as this processed directive.
9 The communication of the content of these and other fields of health and load
10 tables 1204 is accomplished using a message protocol that is described below with
11 reference to FIG. 15.

12 FIG. 15 illustrates an exemplary message protocol 1500 for the health and
13 load information-related communications that are illustrated in FIG. 12 between
14 hosts 108 and load balancing units 106. Generally, an event-driven mechanism is
15 used to push changes to health and load tables 1204 from hosts 108 to load
16 balancing units 106. In other words, for a described implementation, information
17 is transmitted from hosts 108 to load balancing units 106 when health and load
18 tables 1204 are updated. This avoids periodically sending a snapshot of all of each
19 health and load table 1204, which reduces network bandwidth consumption by
20 health and load infrastructure 1202.

21 Message protocol 1500 may be implemented using any available message
22 transport mechanism. Such mechanisms include reliable multicast transmission,
23 point-to-point transmission (e.g., user datagram protocol (UDP)), and so forth. As
24 illustrated, message protocol 1500 includes seven message types 1502-1514: a
25 heartbeat message 1502, a goodbye message 1504, a row change message 1506, a

1 get table snapshot message 1508, a send table snapshot message 1510, a postulate
2 table state message 1512, and a postulate wrong message 1514.

3 It should be understood that, with the exception of arrows 1516 and 1518,
4 no temporal relationship between or among the different messages types 1502-
5 1514 is implied by the illustration. For example, a row change message 1506 does
6 not typically follow a goodbye message 1504.

7 Heartbeat message 1502 indicates that a particular host 108 is functioning
8 and provides some error checking for the content of a corresponding particular
9 health and load table 1204 with respect to a corresponding particular cache for the
10 particular host 1304 in consolidated health and load cache 1208. Each health and
11 load infrastructure 1202 at each host 108 sends a heartbeat message directly or
12 indirectly to each consolidated health and load cache 1208 at each load balancing
13 unit 106.

14 Heartbeat messages 1502 address the aging-out problem for data in
15 consolidated health and load caches 1208 that arises, in part, because a snapshot of
16 the entirety of each health and load table 1204 is not periodically transmitted to
17 each load balancing unit 106. A transmission scheme for heartbeat messages 1502
18 is described further below with reference to FIG. 16.

19 Heartbeat messages 1502 include an identifier for the host, error checking
20 data, and optionally a DNS name. The identifier of the host may be a unique (e.g.,
21 32-bit) number that is selected at configuration time. The error checking data may
22 be a checksum, a state-change sequence number, a generation number, a CRC
23 value, etc. that enables a receiving load balancing unit 106 to validate that the
24 contents of its consolidated health and load cache 1208 comports with the contents
25 of the health and load table 1204 of the transmitting host 108. If a generation

1 number approach is employed, then multiple generation IDs can be used with each
2 generation ID assigned to a “chunk” of applications. Messages can then refer to a
3 chunk number or a chunk number/generation ID pair, depending on the context.

4 The error checking data (or, more generally, a content indicator) may be a
5 single value for the health and load table 1204 overall, or it may be multiple values
6 determined on a per-entry 1302 basis. The DNS name may optionally be sent
7 (e.g., every “x” heartbeats) to verify or update the current correct network address
8 for the host.

9 Goodbye message 1504 is sent from a particular host 108 to load balancing
10 units 106 to indicate that the particular host 108 is planning to shutdown.
11 Goodbye message 1504 includes a host identifier that may be indexed/mapped to a
12 network address for the particular host 108. Goodbye message 1504 is used for
13 clean, intentional shutdowns by hosts 108 to precipitate a “fast clear”. However, if
14 a goodbye message 1504 is lost, caches eventually age out the particular host’s
15 108 entries because heartbeat messages 1502 are no longer sent.

16 Row change message 1506 is sent from a particular host 108 to load
17 balancing units 106 to indicate that the health and/or load for a given application
18 316 of the particular host 108 has changed. Row change message 1506 includes a
19 host identifier, an application identifier, an operation, and data for the operation.
20 Exemplary host identifiers are described above with regard to heartbeat messages
21 1502 and goodbye messages 1504. Exemplary application identifiers are
22 described above with regard to application identifier 1302(A) of an application-
23 associated entry 1302 of health and load tables 1204.

24 The row change operation may be add, delete, or update. In other words,
25 the data for the operation may be added to (for an add operation) or a replacement

1 for (for an update operation) information already present at consolidated health
2 and load caches 1208 at load balancing units 106. For a delete operation, no data
3 need be provided. Message protocol 1500 is defined such that multiple operations
4 may be stipulated to be performed for a single row change message 1506. Hence
5 for a particular host identifier, sets of an application identifier, operation, and
6 operation data may be repeated for multiple applications 316 of the host 108 that is
7 identified by the particular host identifier.

8 Get table snapshot message 1508 is sent from a particular load balancing
9 unit 106 for a particular consolidated health and load cache 1208 to an individual
10 host 108 or hosts 108. This get table snapshot message 1508 requests that health
11 and load infrastructure 1202 at hosts 108 provide a snapshot of the respective
12 health and load table 1204 for the respective host 108. This message includes an
13 identification of the requesting load balancing unit 106 and may be used by a load
14 balancing unit 106 (i) after it has failed and then recovered; (ii) after a host 108
15 fails, recovers, and begins sending heartbeat messages 1502 again; (iii) if a row
16 change message 1506 is sent to load balancing unit 106, but the message gets
17 dropped, so its consolidated health and load cache 1208 is out of sync with the
18 respective health and load table 1204 for the respective host 108; and (iv) so forth.

19 For the third (iii) situation, the lack of synchronization between
20 consolidated health and load cache 1208 and the respective health and load table
21 1204 for the respective host 108 is discovered by a subsequent heartbeat message
22 1502 from the respective host 108 because the “error checking” will indicate that
23 consolidated health and load cache 1208 is out of date. Load balancing unit 106
24 can then send a get table snapshot message 1508 so that it can update its
25 consolidated health and load cache 1208. Thus, for any of the three (i, ii, iii)

1 exemplary situations, load balancing unit 106 subsequently reconstitutes its
2 consolidated health and load cache 1208 using get table snapshot 1508. Get table
3 snapshot 1508 may be sent repeatedly to each host 108 in a point-to-point manner
4 or may be sent one time to many hosts 108 in a multicast manner.

5 Send table snapshot message 1510 is sent from an individual host 108 to a
6 particular load balancing unit 106 after the individual host 108 has received a get
7 table snapshot message 1508 from the particular load balancing unit 106 as
8 indicated by arrow 1516. The contents of a send table snapshot message 1510 is
9 prepared by health and load infrastructure 1202 and may include all or at least
10 multiple rows of the health and load table 1204 of the individual host 108 so that
11 the particular load balancing unit 106 may rebuild its consolidated health and load
12 cache 1208. Send table snapshot message 1510 may be a separately designed
13 message, or it may be equivalent to a sequence of add operations in a row change
14 message 1506.

15 Postulate table state message 1512 and postulate wrong message 1514 are
16 related to the target load balancing state and the current load balancing state of
17 load balancer directive 1302(C) of an entry 1302 in a health and load table 1204.
18 The target load balancing state is the directive that health and load infrastructure
19 1202 desires load balancing units 106 to be operating under. The current load
20 balancing state is the directive that health and load infrastructure 1202 expects or
21 believes that load balancing units 106 are currently operating under. Generally,
22 the two load balancing states are identical.

23 However, the target load balancing state may differ from the current load
24 balancing state during a transitional period for a state directive change. For
25 example, the target load balancing state and the current load balancing state are

1 both initially set to active. A problem with host 108 and/or an application 316
2 thereof is detected and the target load balancing state directive is switched to
3 draining. This draining directive is communicated to load balancing units 106
4 using a row change message 1506.

5 There is a delay before this directive change is noted in all consolidated
6 health and load caches 1208 of all load balancing units 106. During this
7 transitional period, the target load balancing state is draining while the current load
8 balancing state is still active at health and load table 1204 of host 108. Before
9 changing the current load balancing state to draining, health and load
10 infrastructure 1202 wants to ensure that consolidated health and load caches 1208
11 have actually been updated to the new directive state of draining.

12 To verify that consolidated health and load caches 1208 of load balancing
13 units 106 have been updated to a new state directive, health and load infrastructure
14 1202 sends a postulate table state message 1512 to load balancing units 106.
15 Postulate table state message 1512 is sent some time (e.g., a predetermined delay
16 period) after transmission of a row change message 1506 indicating that the state
17 directive is to be changed. The postulate table state message 1512, in this
18 example, indicates that the table state should be draining. As indicated by the
19 dashed arrow 1518, a load balancing unit 106 responds to this postulate table state
20 message 1512 if its consolidated health and load cache 1208 differs from the
21 postulated state directive.

22 If the directive in consolidated health and load cache 1208 does differ from
23 the postulated state directive, then that load balancing unit 106 sends a postulate
24 wrong message 1514 to the health and load infrastructure 1202 of the host 108 that
25 issued the postulate table state message 1512. This health and load infrastructure

1202 then periodically resends postulate table state message 1512 until no further
postulate wrong messages 1514 are received from consolidated health and load
caches 1208. At that point, health and load infrastructure 1202 sends a row change
message 1506 with the new current load balancing state. In this sense,
consolidated health and load caches 1208 are the definitive determiners of the
current load balancing state, and health and load infrastructure 1202 is the
definitive determiner of the target load balancing state.

FIG. 16 illustrates an exemplary message transmission scheme for the
communications that are illustrated in FIG. 12 between hosts 108 and load
balancing units 106. The exemplary message transmission scheme can reduce the
bandwidth consumed by heartbeat messages 1502 on communication linkage
1210. The message transmission scheme of FIG. 16 is particularly adapted to
heartbeat messages 1502, but it may also be utilized for other messages of
message protocol 1500.

A group of hosts 108(1), 108(2), 108(3) ... 108(11), and 108(12) are
illustrated along with load balancing units 106(1), 106(2) ... 106(u). Each line
represents membership linkage or inclusion among the group of hosts 108(1, 2 ...
12). The group of hosts 108(1, 2 ... 12) form a membership of nodes that work
together to propagate heartbeat information to load balancing units 106. Although
twelve hosts are shown, more or fewer may be part of any given group of hosts.
Also, a total set of hosts 108 that are being served by a load balancing
infrastructure 106 may be divided into one, two, three, or more groups of hosts.

In a described implementation, the membership of nodes for group of hosts
108(1, 2 ... 12) elect a leader that is responsible for transmitting heartbeat
messages 1502 to load balancing units 106. Each (non-leading) host 108 in group

1 of hosts 108(1, 2 ... 12) sends its heartbeat messages 1502 to the elected leader.

2 Host 108(4) is the elected leader in this example.

3 With the membership of nodes, heartbeat information for each host 108 in
4 group of hosts 108(1, 2 ... 12) propagates to the group leader host 108(4). Host
5 108(4) collects the heartbeat information and consolidates it into a consolidated
6 heartbeat message 1602. Consolidated heartbeat messages 1602(1), 1602(2) ...
7 1602(u) are then sent to respective load balancing units 106(1), 106(2) ... 106(u).
8 These consolidated heartbeat messages 1602 may optionally be compressed to
9 further reduce bandwidth consumption.

10 As another alternative, the leader host 108(4) may only forward changes in
11 group membership to consolidated health and load caches 1208. In other words,
12 in this mode, consolidated health and load caches 1208 deal primarily if not solely
13 with state changes to membership. It is the responsibility of the leader host 108(4)
14 to ensure that the first hello is forwarded when a host 108 comes online and that a
15 goodbye message 1504 gets sent when that host 108 goes offline. Additionally, a
16 host 108 can periodically specify that a heartbeat message 1502 is to be
17 "forwarded". This indicates to the leader host 108(4) to send it to consolidated
18 health and load caches 1208 even though it does not represent a membership
19 change.

20 Heartbeat messages 1502 (including consolidated heartbeat messages 1602)
21 are used by load balancing units 106 when their consolidated health and load
22 caches 1208 are unsynchronized with health and load tables 1204. This lack of
23 synchronization may arise, for example, from a crash or other failure of
24 consolidated health and load cache 1208 and/or of load balancing unit 106. As
25 described above, each heartbeat message 1502 includes error checking data that is

1 usable to verify equivalency between a consolidated health and load cache 1208
2 and health and load tables 1204. If non-equivalency is discovered with regard to a
3 particular host 108 and/or an application 316 thereof, the DNS name of the
4 particular host 108 is acquired from the heartbeat messages 1502.

5 The DNS name is used by consolidated health and load cache 1208 to send
6 a get table snapshot message 1508 to the particular host 108 in order to get
7 updated health and load information 1206 in the form of a send table snapshot
8 message 1510. A different or the same get table snapshot message 1508 is sent to
9 each host 108 for which non-equivalency is discovered. Eventually, the health and
10 load information 1206 in the consolidated health and load cache 1208 is equivalent
11 to the health and load information 1206 in health and load tables 1204 as verifiable
12 by new heartbeat messages 1502. In this manner, a failed consolidated health and
13 load cache 1208 can be bootstrapped back into operation without manual oversight
14 using message protocol 1500 and an equivalency-checking scheme.

15 FIG. 17A and FIG. 17B illustrate exemplary health and load information
16 proxy storage scenarios for health and load tables 1204 and for consolidated health
17 and load caches 1208, respectively. In implementations described above with
18 reference to FIGS. 12-16, hosts 108 include health and load infrastructure 1202.
19 However, other implementations may entail hosts that do not include health and
20 load infrastructure 1202.

21 For example, a host may be running a version of application(s) and/or an
22 operating system for which health and load infrastructure is either not
23 implemented or for policy reasons may not be installed on the host. Consequently,
24 these types of hosts do not have health and load infrastructure 1202 executing
25 thereon. Host 1702 is such a host that does not execute health and load

1 infrastructure 1202. Nevertheless, host 1702 can utilize a health and load
2 infrastructure 1202 that is executing on one or more proxies, such as proxy 1704.

3 Proxy 1704 has resident thereat and executing thereon a health and load
4 infrastructure 1202, which includes a health and load table 1204. Host 1702 can
5 use the functionality of health and load infrastructure 1202 by providing health
6 and load information 1206 to health and load table 1204 for applications that are
7 running on host 1702. Alternatively, proxy 1704 can deduce health and load on
8 host 1702 by performing external monitoring actions. Proxy 1704 is illustrated as
9 proxy 1704(1) and 1704(2) for redundancy and the resulting high availability.

10 In implementations described above with reference to FIGS. 12-16 and
11 below with reference to FIG. 18, load balancing is effectuated with load balancing
12 units 106 that include consolidated health and load caches 1208. However, other
13 implementations may entail load balancing that does not include consolidated
14 health and load caches 1208.

15 For example, load balancing may be effectuated by monolithic load
16 balancing hardware or other load balancing infrastructure that does not and/or
17 cannot store or otherwise include a consolidated health and load cache 1208. Load
18 balancer 1706 reflects such a load balancing device or devices that do not have a
19 consolidated health and load cache 1208. Nevertheless, load balancer 1706 can
20 utilize a consolidated health and load cache 1208 that exists on one or more
21 proxies, such as proxy 1708.

22 Proxy 1708 includes a consolidated health and load cache 1208, which
23 stores health and load information 1206 for hosted applications being serviced by
24 load balancer 1706. Load balancer 1706 can use the health and load information
25 1206 of consolidated health and load cache 1208 when performing load balancing

functions by accessing such information using application programming interfaces (APIs) native to and supported by load balancer 1706. Alternatively, consolidated health and load cache 1208 can invoke APIs to push health and load information 1206, including directives, to load balancer 1706. Proxy 1708 is illustrated as proxy 1708(1) and 1708(2) for redundancy and the resulting high availability.

FIG. 18 illustrates an exemplary target application endpoint allotment procedure that involves a classifier 304 and a health and load handler 314 of a load balancing unit 106. After health and load handler 314 has acquired a consolidated health and load cache 1208, health and load information 1206 thereof is utilized in the selection of application endpoints for new requests/connections.

As described above with reference to FIG. 13B, consolidated health and load cache 1208 includes cached health and load information 1206 for multiple hosts 108. To facilitate the creation and updating of consolidated health and load cache 1208 from health and load information 1206 that originates from multiple hosts 108, the health and load information 1206 therein is organized so that it may be accessed by identifier of each host 108. However, the health and load information 1206 therein is also organized such that it can be accessed by type of application 316 in order to facilitate application endpoint selection.

In other words, health and load handler 314 is capable of accessing health and load information 1206 on a per-application 316 basis across health and load information 1206 for multiple hosts 108. Once health and load information 1206 for a given application 316 has been accessed for each host 108, allocation of incoming connection requests may be performed in accordance with such health and load information 1206. For example, possible endpoints for the given application 316 may be allocated to incoming connection requests by selection of

1 the endpoints of the given application 316 with consideration of available relative
2 load capacity among healthy endpoints for the given application 316.

3 In a described implementation, classifier 304 makes a target application
4 endpoint allotment request 1802 to health and load handler 314. As illustrated,
5 target application endpoint allotment request 1802 includes (i) a virtual IP address
6 and port, (ii) a protocol, and (iii) protocol-specification information. Target
7 application endpoint allotment request 1802 therefore identifies a type of
8 application 316 to which incoming connection requests are directed.

9 Health and load handler 314 receives target application endpoint allotment
10 request 1802 and selects at least one physical endpoint corresponding to the
11 identified type of application 316 using any one or more of many selection
12 mechanisms. To reduce latency, health and load handler 314 selects an allotment
13 of application endpoints to be used over a number of incoming connection
14 requests. This allotment is provided from health and load handler 314 to classifier
15 304 using target application endpoint allotment response 1804. As illustrated,
16 target application endpoint allotment response 1804 includes an allotment of
17 physical IP addresses and ports (such as endpoints IP1, IP2, and IP3) for the
18 identified type of application 316.

19 The allotment for target application endpoint allotment response 1804 may
20 be completed using one or more allotment schemes. By way of example, a token
21 allotment scheme 1806 and a percentage allotment scheme 1808 are illustrated.
22 Token allotment scheme 1806 is a unit-based allotment scheme, and percentage
23 allotment scheme 1808 is a time-based allotment scheme.

24 Token allotment scheme 1806 allocates tokens for each healthy endpoint
25 IP1, IP2, and IP3 responsive to their relative load and capacity ratios. For the

example as illustrated, of the total available capacity, IP1 has 40% of the available capacity, IP2 has 35% of the available capacity, and IP3 has 25% of the available capacity. Thus, the total number of tokens is divided along these percentages. The total number of tokens may be provided as part of target application endpoint allotment request 1802 or determined by health and load handler 314.

Any value for the total number of tokens may be used, such as 10, 45, 100, 250, 637, 1000, and so forth. This value may be set in dependence on the number of connection requests per second and the speed/frequency at which application health and/or load is changing. Classifier 304 “uses up”/consumes one token when responding to each connection request with an application endpoint allocation until the tokens are exhausted; classifier 304 then requests another token allotment using target application endpoint allotment request 1802.

Percentage allotment scheme 1808 determines available relative capacity in a similar manner. However, instead of tokens, these determined available relative capacities per application endpoint are provided to classifier 304 along with a duration timer 1810. Classifier 304 allocates target application endpoints to incoming connection requests in accordance with these available relative capacity percentages until expiration of duration timer 1810.

For percentage allotment scheme 1808, classifier 304 maintains a running record of application endpoint allocations to adhere to the allotted percentages and keeps track of time for duration timer 1810. When the timer expires, classifier 304 then requests another percentage allotment using target application endpoint allotment request 1802.

It should be noted that token allotment scheme 1806 can also use a time limit. If allotted tokens are too old, they should be discarded and new ones

1 acquired. Otherwise, classifier 304 may consume stale tokens that were
2 previously allocated based on health and load information that is currently too
3 outdated. Use of application endpoint allotments by classifier 304 is described
4 further below in the section entitled “Exemplary Classifying, Forwarding, and
5 Request Routing”.

6 **Exemplary Session Tracking**

7 This section describes how host status information, such as session
8 information, may be collected for and utilized in network load balancing. This
9 section primarily references FIGS. 19-24 and illuminates session affinity
10 preservation functionality such as that provided by session tracker 308 (of FIG. 3).
11 As described above with reference to FIGS. 1-3, each host 108 hosts one or more
12 applications 316 that provide service(s) to clients 102. Session tracker 308 utilizes
13 session information that relates to contexts for the connections established
14 between applications 316 and clients 102 for certain described implementations of
15 network load balancing.

16 FIG. 19 illustrates an exemplary network load balancing approach that
17 involves session information 1902. At connection [1], client 102(1) is shown
18 making a new connection with host 108(2) via load balancing infrastructure 106.
19 Load balancing infrastructure 106 may be comprised of one or more load
20 balancing units 106. When the connection request arrives at load balancing
21 infrastructure 106, the request is typically routed to a host 108 using network load
22 balancing functionality responsive to health and/or load information of hosts 108
23 and/or applications 316 (not explicitly shown in FIG. 19) thereof.

24 When connection [1] is made, a session is established between client 102(1)
25 and the servicing application 316, which is on host 108(2) in this example. The

1 session provides a context for the communication exchange between client 102(1)
2 and host 108(2). The information for the session context is stored at host 108(2).
3 When connection [1] is completed, the session context may not be used again. On
4 the other hand, the session context may be useful again if client 102(1) attempts to
5 initiate another connection with hosts 108 for the service provided by application
6 316. If this other connection is not routed to the same host 108(2) that stores that
7 session context, then client 102(1) has to establish a new session context, which
8 can be time consuming, data/processing intensive, and/or frustrating to the human
9 user of client 102(1). With health and/or load information-based network load
10 balancing, there is no likelihood greater than random chance that the second
11 connection will be routed to 108(2).

12 However, if load balancing infrastructure 106 has access to a mapping
13 between session information and hosts 108, load balancing infrastructure 106 can
14 route connection requests that relate to previously established sessions to the
15 appropriate host 108. Some session information may be inferred from the contents
16 of packets flowing through load balancing infrastructure 106. However, this
17 approach is imprecise and haphazard for a number of reasons. First, session
18 establishment and termination is merely inferred. Second, some sessions are not
19 “officially” terminated with an appropriate indication that is included in a packet.
20 For example, some sessions simply time out. Third, packets being transmitted
21 from host 108(2) to client 102(1) may take a path that does not include load
22 balancing infrastructure 106, which precludes any snooping of such packets by
23 load balancing infrastructure 106 for session information.

24 As shown in FIG. 19, hosts 108 provide session information (SI) 1902 to
25 load balancing infrastructure 106. Using session information 1902 from hosts

1 108, a session affinity preserver 1904 can preserve the affinity between an
2 established session and the host 108 on which the session was established.
3 Session information 1902 includes a linkage between or a mapping from each
4 session established between a client 102 and a particular host 108 to that particular
5 host 108. This mapping is accessible to session affinity preserver 1904 as part of
6 host-session information mapping 1906. More-specific examples of session
7 information 1902 are provided below especially with reference to FIGS. 20, 22,
8 23A, and 23B.

9 In certain described implementations for session tracking, the logical nature
10 of clients 102 is pertinent. As noted above with reference to FIG. 1, a client 102
11 may be a specific device and/or a specific user of a device. Consequently, session
12 affinity for a user client 102 that is accessing hosts 108 from different devices can
13 still be preserved. Session continuations using session information 1902 can
14 therefore still be effectuated in proxy scenarios (e.g., those of some internet
15 service providers (ISPs)).

16 Continuing with the connection [1] example, the session established at host
17 108(2) is provided to load balancing infrastructure 106 as session information
18 1902. Specifically, a linkage/mapping between (i) the session context of client
19 102(1) and host 108(2) and (ii) an identifier for host 108(2) is created at host-
20 session information mapping 1906. When a connection request for connection [2]
21 subsequently arrives for the same session context, session affinity preserver 1904
22 locates this session context in host-session information mapping 1906 and
23 ascertains that host 108(2) is associated with this session context from the
24 linkage/mapping.

1 Responsive to the mapping of host 108(2) to the requested session context
2 as ascertained by session affinity preserver 1904 from host-session information
3 mapping 1906, connection [2] is routed to host 108(2). In this sense, preserving
4 session affinity is a higher priority for load balancing infrastructure 106 than
5 application health and load-based network load balancing decisions. However,
6 health and/or load may be a more important network load balancing factor than
7 session tracking when, for example, loading is extremely heavy or when the
8 session-relevant application and/or host is in a failed condition.

9 Many types of connections may be session-related. Examples include: a
10 TCP connection, a transport layer security (TLS)/SSL session, a PPTP session, an
11 IPSec/L2TP session, an ISA session, an HTTP cookie-based session, a Terminal
12 Server session, an administrator-defined session, and so forth. By way of
13 clarification, a TCP connection is considered to be a session of TCP packets. Also,
14 a model for defining sessions by an administrator may be enumerated and
15 supported. Furthermore, client IP-address-based sessions that are delineated by
16 timeouts may also be supported. This is relatively non-intelligent session support,
17 but is expected by some users.

18 A connection request from a client 102 varies by the type of desired
19 session. For example, for sessions of type “TCP connection”, the connection
20 request comprises a TCP packet. For sessions of type “SSL session”, the
21 connection request comprises a TCP connection. Other such connection requests
22 correspond to other session types. These examples also show how there may be
23 session layers. At a lower session level, a session context for a TCP connection
24 may include a TCP 4-tuple, a session number, the number of bytes sent/received,
25 and so forth. At a higher session level, a session context for an SSL session may

1 include a 32-byte session ID, a public key of the client 102 that is provided to the
2 host 108, and so forth.

3 FIG. 20 illustrates an exemplary network load balancing approach that
4 involves communicating session information using notifications 2006 and
5 messages 2008. Multiple load balancing units 106(1), 106(2) ... 106(u) and
6 multiple hosts 108(1), 108(2) ... 108(n) are shown. Each respective host 108(1),
7 108(2) ... 108(n) includes one or more respective applications 316(1), 316(2) ...
8 316(n) which are resident thereat and executing thereon. Notifications 2006 are
9 used to provide session information from applications 316, and messages 2008 are
10 used to provide session information from hosts 108 to load balancing units 106.

11 As illustrated, each respective host 108(1), 108(2) ... 108(n) includes
12 respective session tracking infrastructure (STI) 2002(1), 2002(2) ... 2002(n).
13 Each respective session tracking infrastructure 2002(1), 2002(2) ... 2002(n)
14 includes a respective session table 2014(1), 2014(2) ... 2014(n) (although only
15 session table 2014(1) is explicitly illustrated in FIG. 19).

16 Each respective load balancing unit 106(1), 106(2) ... 106(u) includes
17 respective traffic routing functionality (TRF) 2012(1), 2012(2) ... 2012(u). Traffic
18 routing functionality 2012 may comprise, for example, classifying and/or
19 requesting routing functionality, such as that provided by classifier 304 and
20 request router 306, respectively. Distributed across load balancing units 106(1),
21 106(2) ... 106(u) is a distributed session tracking manager 2010.

22 In a described implementation, traffic routing functionality 2012 and
23 distributed session tracking manager 2010 are part of load balancing infrastructure
24 106. Session tracking infrastructure 2002 may also be (e.g., a remote) part of load
25 balancing infrastructure 106.

1 An API 2004 is employed to provide session information from applications
2 316 to session tracking infrastructure 2002. Using API 2004, applications 316 are
3 empowered to notify session tracking infrastructure 2002 of session information,
4 including various changes thereto. More specifically, each application 316 is
5 capable of providing, and session tracking infrastructure 2002 is capable of
6 accepting, notifications 2006.

7 A notification that a session has been established (or session establishment
8 notification 2006(E)) is provided from application 316 when a session is newly
9 established or opened. Session establishment notification 2006(E) includes a
10 session identifier and optionally an identifier of application 316. A notification
11 that a session has been terminated (or session termination notification 2006(T)) is
12 provided from application 316 when a session is terminated or closed. Session
13 termination notification 2006(T) also includes the session identifier and optionally
14 the identifier of application 316.

15 When session tracking infrastructure 2002 accepts a session establishment
16 notification 2006(E), it inserts an entry in session table 2014 for the new session.
17 An exemplary session table 2014 is described further below with reference to FIG.
18 23A. When session tracking infrastructure 2002 accepts a session termination
19 notification 2006(T), it removes the entry in session table 2014 for the old session.

20 Session table 2014(1) is the authoritative source for session information
1902 with respect to applications 316(1) on host 108(1). There is generally too
21 much latency, however, to require traffic routing functionality 2012 to contact
22 hosts 108 for access to session tables 2014 upon receipt of each incoming
23 connection request having a session reference. Session information 1902 is
24 therefore cached at load balancing units 106.

1 At load balancing units 106, distributed session tracking manager 2010
2 caches session information 1902 as part of its session tracking management
3 responsibilities. Generally, distributed session tracking manager 2010 is a
4 distributed application and/or virtual service that resides partially on each load
5 balancing unit 106. For each logical session, distributed session tracking manager
6 2010 keeps at least one cached copy of session information therefor in a reliable
7 and scalable manner that may be quickly utilized for routing traffic as incoming
8 connection requests that have a session reference are received by load balancing
9 infrastructure 106.

10 Communications between hosts 108 and load balancing units 106 are
11 effectuated with a reliable protocol that ensures that messages 2008 sent from a
12 host 108 arrive at the intended load balancing unit 106. Each host 108 is bound to
13 at least one specific load balancing unit 106 that is the intended load balancing
14 unit 106 for messages 2008. This binding is created by assigning an IP address of
15 a specific load balancing unit 106 to each host 108 for sending session-tracking
16 messages 2008 between session tracking infrastructure 2002 and distributed
17 session tracking manager 2010. To facilitate high availability of load balancing
18 infrastructure 106, if a load balancing unit 106 fails, another load balancing unit
19 106 assumes the IP address of the failed load balancing unit 106. Failure detection
20 for IP address assumption may be accomplished using a heartbeat or another
21 aliveness monitoring scheme.

22 Thus, messages 2008 communicate session information 1902 from session
23 tracking infrastructure 2002 to distributed session tracking manager 2010. For
24 example, when session tracking infrastructure 2002 accepts a session
25 establishment notification 2006(E), it also sends a session up message 2008(U) to

1 distributed session tracking manager 2010. Session up message 2008(U) includes
2 the session identifier, a host identifier, and optionally other information. Contents
3 for a session up message 2008(U) are described further below with reference to
4 FIG. 23B with respect to information that may be stored for each session by an
5 implementation of distributed session tracking manager 2010. When session
6 tracking infrastructure 2002 accepts a session termination notification 2006(T), it
7 also sends a session down message 2008(D) to distributed session tracking
8 manager 2010. Messages 2008 can be sent before, during, or after session
9 tracking infrastructure 2002 appropriately modifies session table 2014 in response
10 to notifications 2006.

11 FIG. 21 is a flow diagram 2100 that illustrates an exemplary method for
12 network load balancing that involves communicating session information using
13 notifications and messages. Flow diagram 2100 includes fifteen blocks 2102-
14 2130. Although the actions of flow diagram 2100 may be performed in other
15 environments and with a variety of software schemes, FIGS. 1-3 and 19-20 are
16 used in particular to illustrate certain aspects and examples of the method.

17 For example, the actions of four blocks 2102-2104 and 2118-2120 are
18 performed by an application 316, the actions of six blocks 2106-2110 and 2122-
19 2126 are performed by session tracking infrastructure 2002, and the actions of five
20 blocks 2112-2116 and 2128-2130 are performed by a distributed session tracking
21 manager 2010. The actions of eight of these blocks 2102-2116 are primarily
22 directed to opening a session, and the actions of seven of these blocks 2118-2130
23 are primarily directed to closing a session.

24 At block 2102, a session is opened. For example, application 316 may
25 open a session with a client 102. At block 2104, a session establishment

1 notification is provided. For example, application 316 may provide a session
2 establishment notification 2006(E) to session tracking infrastructure 2002 using
3 API 2004 as a consequence of and/or in conjunction with opening the session.

4 At block 2106, the session establishment notification is accepted. For
5 example, session tracking infrastructure 2002 may accept session establishment
6 notification 2006(E) from application 316 in accordance with API 2004. At block
7 2108, an entry in a session table is inserted. For example, session tracking
8 infrastructure 2002 may insert an entry in session table 2014 for the opened
9 session. Examples of such insertion are described further below especially with
10 reference to FIG. 23A. At block 2110, a session up message is sent. For example,
11 session tracking infrastructure 2002 may send a session up message 2008(U) to
12 distributed session tracking manager 2010 using a reliable communication
13 protocol.

14 At block 2112, the session up message is received. For example,
15 distributed session tracking manager 2010 may receive session up message
16 2008(U) from session tracking infrastructure 2002 in accordance with the reliable
17 communication protocol. At block 2114, a session information entry is created.
18 For example, distributed session tracking manager 2010 may create a session
19 information entry for cached session information 1902 at one or more load
20 balancing units 106. Examples of such creating and subsequent adding are
21 described further below especially with reference to FIGS. 22 and 23B.

22 At block 2116, network traffic is routed with the session information. For
23 example, traffic routing functionality 2012 in conjunction with distributed session
24 tracking manager 2010 may use cached session information 1902, including the
25 created session information entry, to route incoming connection requests that have

1 a session reference. An example of such traffic routing is described further below
2 especially with reference to FIG. 24. Additional examples are described below in
3 the section entitled “Exemplary Classifying, Forwarding, and Request Routing”.

4 At block 2118, the session is closed. For example, application 316 may
5 close the session with client 102. At block 2120, a session termination notification
6 is provided. For example, application 316 may provide a session termination
7 notification 2006(T) to session tracking infrastructure 2002 using API 2004 as a
8 consequence of and/or in conjunction with closing the session.

9 At block 2122, the session termination notification is accepted. For
10 example, session tracking infrastructure 2002 may accept session termination
11 notification 2006(T) from application 316 in accordance with API 2004. At block
12 2124, the entry in the session table is removed. For example, session tracking
13 infrastructure 2002 may remove the entry in session table 2014 for the closed
14 session. At block 2126, a session down message is sent. For example, session
15 tracking infrastructure 2002 may send a session down message 2008(D) to
16 distributed session tracking manager 2010 using the reliable communication
17 protocol.

18 At block 2128, the session down message is received. For example,
19 distributed session tracking manager 2010 may receive session down message
20 2008(D) from session tracking infrastructure 2002 in accordance with the reliable
21 communication protocol. At block 2130, the session information entry is
22 destroyed. For example, distributed session tracking manager 2010 may destroy
23 the session information entry at the cached session information 1902 at any load
24 balancing units 106 that have the session information entry. Examples of such
25

1 destroying and subsequent deleting are described further below especially with
2 reference to FIGS. 22 and 23B.

3 FIG. 22 illustrates an exemplary approach to managing session information
4 at multiple load balancing units 106. Each respective load balancing unit 106(1),
5 106(2) ... 106(u) includes a respective part 2202(1), 2202(2) ... 2202(u) of a
6 distributed atom manager (DAM) 2202. DAM 2202 is an exemplary
7 implementation of distributed session tracking manager 2010. Each respective
8 DAM portion 2202(1), 2202(2) ... 2202(u) includes a respective part 2206(1),
9 2206(2) ... 2206(u) of a DAM table (DAMT) 2206.

10 DAM 2202 is a distributed application or virtual service that manages
11 session information 1902 in a reliable and scalable manner so that traffic routing
12 functionality 2012 can use it to preserve session affinity. For example, traffic
13 routing functionality 2012 can access DAM 2202 using an API (not specifically
14 shown) to search or have searched DAMT 2206. Function calls 2204, operation of
15 DAM 2202, and other aspects of FIG. 22 are described further below after the
16 description of FIGS. 23A and 23B.

17 FIG. 23A is an exemplary session table 2014 as illustrated in FIG. 20.
18 Session table 2014 includes “v” entries 2302(1), 2302(2) ... 2302(v). Each entry
19 2302 is inserted by session tracking infrastructure 2002 responsive to a session
20 establishment notification 2006(E) that is accepted from an application 316. Each
21 entry 2302 is removed by session tracking infrastructure 2002 responsive to a
22 session termination notification 2006(T) that is accepted from application 316.

23 As described above, each session establishment notification 2006(E)
24 includes a session identifier and optionally an identifier of application 316. Each
25 respective entry 2302(1), 2302(2) ... 2302(v) in session table 2014 includes

1 respective fields of (i) session identifier 2302(1I), 2302(2I) ... 2302(vI) and (ii)
2 session type and/or application 2302(1T), 2302(2T) ... 2302(vT).

3 Session type and/or application 2302(T) may be “TCP”, “IPSEC”,
4 “Terminal Server,” “HTTP-cookie”, an application type as noted above, and so
5 forth. Session identifier 2302(I) may be “<source IP address, source TCP port,
6 destination IP address, destination TCP port >”, “Client IP = 172.30.189.122”,
7 “User = ‘joe_user’”, “Cookie = ‘{b7595cc9-e68b-4eb0-9bf1-bb717b31d447}’”,
8 another e.g. application-specific identification for a session, and so forth. For TCP
9 connection/session types, session identifier 2302(I) may alternatively be a
10 canonical version of the TCP 4-tuple (for IPv4 or IPv6). Other values for the
11 fields of session identifier 2302(I) and application/session type 2302(T) may
12 alternatively be used.

13 FIG. 23B is an exemplary distributed atom manager (DAM) table (DAMT)
14 2206 as illustrated in FIG. 22. DAM table 2206 includes “w” entries 2304(1),
15 2304(2) ... 2304(w). Each session information entry 2304 is created by DAM
16 2202 responsive to a session up message 2008(U) that is received from session
17 tracking infrastructure 2002. Each session information entry 2304 is destroyed
18 responsive to a session down message 2008(D) that is received from session
19 tracking infrastructure 2002. As described further below, session information
20 entries 2304 of DAM tables 2206 may actually be manipulated by DAM 2202
21 using function calls 2204.

22 As described above, session up message 2008(U) includes the session
23 identifier, a host identifier, and optionally other information. Each respective
24 session information entry 2304(1), 2304(2) ... 2304(w) in DAM table 2206
25 includes respective fields of (i) key 2304(1K), 2304(2K) ... 2304(wK), (ii) data

1 2304(1D), 2304(2D) ... 2304(wD), and (iii) metadata 2304(1M), 2304(2M) ...
2 2304(wM). For example, values for key 2304(K) fields may be alphanumeric
3 strings, and values for data 2304(D) fields may be binary bits. Values for key
4 2304(K) may be binary bits, too.

5 Key 2304(K) may correspond to the session identifier 2302(I). Data
6 2304(D) may correspond to the host identifier, such as a network address of the
7 host 108 on which the session context exists. Metadata 2304(M) may correspond
8 to other, optional information. Examples of such metadata 2304(M) include data
9 that is used internally by DAM 2202 to resolve atom collisions and to track atom
10 aliveness (e.g., via a time-out mechanism). (This characterization of entries 2304
11 as being atomic is described more fully in the following paragraph.) More
12 specifically, metadata 2304(M) includes, among other things, the identity of the
13 entity (e.g., the instance of traffic routing functionality 2012) that added the
14 session information entry 2304 to the DAM table 2206.

15 In a described implementation, each session information entry 2304 is
16 atomic in the sense that DAM 2202 may add, delete, copy, etc. the entries 2304 as
17 a whole, but DAM 2202 does not ordinarily modify a portion of any whole entry
18 2304. Thus, atomic entries 2304 are added, deleted, copied, otherwise
19 manipulated, etc. across DAM tables 2206 by DAM 2202 in order to implement
20 availability and scalability for a session affinity preservation implementation.

21 Function calls 2204 (of FIG. 22) are usable by DAM 2202 to manipulate the
22 atomic entries 2304 of DAM table 2206. Function calls 2204 may be
23 communicated from one load balancing unit 106 to one or more other load
24 balancing units 106 in a point-to-point or a multicast manner. These function calls
25

1 include add atom 2204(A), delete atom 2204(D), query atom 2204(Q), and return
2 atom 2204(R).

3 Add atom 2204(A) takes the form AddAtom(key, data) and is used to add
4 an atomic entry 2304 to one or more DAM tables 2206. Hence, an add atom
5 2204(A) function call may be formulated as AddAtom(<session identifier>, host
6 IP address). Delete atom 2204(D) takes the form DeleteAtom(key) and is used to
7 delete an atomic entry 2304 at one or more DAM tables 2206. Delete atom
8 2204(D) function calls may be directed at those DAM tables 2206 known to have
9 a copy of the session that is identified by the key 2304(K) or may be multicast to
10 all DAM tables 2206 to ensure that any copies are deleted.

11 Query atom 2204(Q) takes the form QueryAtom(key) and is used by a
12 particular DAM portion 2202 when a session identifier as referenced by an
13 incoming connection request is not located in the particular local DAM table 2206
14 of the particular DAM portion 2202. Query atom 2204(Q) function calls are sent
15 to one or more (including possibly all) other DAM portions 2202. In response,
16 each other DAM portion 2202 checks its local DAM table 2206 for the
17 key/session identifier. If the key is located by another DAM portion 2202, this
18 other DAM portion 2202 replies with a return atom 2204(R).

19 Return atom 2204(R) takes the form ReturnAtom(key, data) and is used to
20 reply to a query atom 2204(Q) function call. Return atom 2204(R) function calls
21 are used when a DAM portion 2202 has a requested atomic entry 2304 in its local
22 DAM table 2206 as identified by a key 2304(K) specified in the query atom
23 2204(Q) function call. Return atom 2204(R) function calls may be directed back
24 to the DAM portion 2202 that issued the query atom 2204(Q) function call.
25

1 Add atom 2204(A) function calls are used in response to session up
2 messages 2008(U) and/or to replicate an atomic entry 2304 to one or more other
3 DAM tables 2206. Such replication may be for redundancy and/or scalability.

4 Delete atom 2204(D) function calls are used in response to session down
5 messages 2008(D) and may also be sent to one or more other DAM tables 2206.
6 After an atomic entry 2304 is deleted, the atomic entry 2304 may enter a “zombie”
7 state such that it remains with DAM 2202, and optionally so that it is actually still
8 stored with DAM table 2206 with a zombie indication in the metadata 2304(M)
9 field of the atomic entry 2304.

10 Thus, once an atomic entry 2304 is deleted, it may stay on in DAM 2202
11 and DAM table 2206 in a zombie state so that packets for this (now dead and
12 closed) session are directed to the host 108 of the session context for proper,
13 protocol-specific treatment. For example, TCP packets received after a TCP
14 connection has been torn down are directed to the host 108 that terminated the
15 connection. This host 108 can respond appropriately – perhaps by sending an RST
16 or by resending a FIN-ACK. The time the atomic entry 2304 spends in this
17 zombie state matches (as closely as reasonably possible) the protocol-specific dead
18 time of the reliable communication protocol that is employed.

19 A query atom 2204(Q) function call is used to attain an atomic entry 2304
20 when a first load balancing unit 106 receives an incoming connection request that
21 references a session that is not stored in the local DAM table 2206 of the DAM
22 2202 of the first load balancing unit 106. It should be noted that other DAM
23 portions 2202 may be queried simultaneously in a broadcast query atom 2204(Q)
24 function call or sequentially until a positive return atom 2204(R) function call is
25 received.

1 A return atom 2204(R) function call is used by a DAM portion 2202 of a
2 second load balancing unit 106 to provide an atomic entry 2304 to the DAM
3 portion 2202 of the first load balancing unit 106, where the atomic entry 2304 has
4 a key 2304(K) that is specified by the key/session identifier in a query atom
5 2204(Q) function call, which was previously issued by the DAM portion 2202 of
6 the first load balancing unit 106. It should be noted that other components, such
7 as traffic routing functionality 2012, may also be capable of calling functions
8 2204, especially a query atom 2204(Q) function call, in accordance with an API or
9 similar.

10 DAM portions 2202 and DAM tables 2206 may be organized and managed
11 in a myriad of manners. Exemplary manners relate to replication/redundancy,
12 local caching upon acquisition, hashing for location selection, and so forth. Zero,
13 one, two, or more levels of replication up to full replication may be employed.
14 With a zero level of replication, each atomic entry 2304 is stored at the DAM 2202
15 that receives a session up message 2008(U) therefor without replication to other
16 DAM portions 2202.

17 With a first level of replication, each atomic entry 2304 is stored at the
18 DAM 2202 that receives a session up message 2008(U) therefor, and it is also
19 added (copied) to one other DAM portion 2202 using an add atom 2204(A)
20 function call. This handles one level of failure for a load balancing unit 106.
21 Similarly, with a second level of replication, each atomic entry 2304 is stored at
22 the DAM 2202 that receives a session up message 2008(U) therefor, and it is also
23 added to two other DAM portions 2202. Generally, the one, two, etc. other DAM
24 portions 2202 to which a given DAM portion 2202 copies atomic entries 2304 is

1 predetermined or selected at random. Third, fourth, etc. levels of replication may
2 also be employed.

3 Furthermore, full replication may be employed by having each atomic entry
4 2304 that is stored at the DAM 2202 that receives a session up message 2008(U)
5 therefor also being added to every other DAM portion 2202. Several factors are
6 impacted by selection of the replication level: As the replication level increases,
7 availability increases and latency decreases. On the other hand, network traffic
8 and memory usage both increase as the replication level increases.

9 When full replication is not employed, local caching upon acquisition may
10 be. For example, when a DAM portion 2202 does not locate a referenced session
11 identifier in its part of DAM table 2206, the DAM portion 2202 issues a query
12 atom 2204(Q) function call to attain the atomic entry 2304 associated with the
13 referenced session identifier via a return atom 2204(R) function call. Instead of
14 jettisoning the attained atomic entry 2304 after use thereof, the DAM portion 2202
15 caches the attained atomic entry 2304 in its part of DAM table 2206. This option
16 offers a tradeoff between the above-enumerated factors.

17 As another option when full replication is not employed, hashing for
18 location selection may be. The first atomic entry 2304 for a session is stored at the
19 DAM portion 2202 that receives the session up message 2008(U). Replicated
20 copy or copies are sent via add atom 2204(A) function calls to specific DAM
21 portion(s) 2202 using a hashing function. Of a total range of possible hash values,
22 each DAM portion 2202 is assigned a subset thereof. Each session identifier is
23 hashed using some hashing function to arrive at a hashing value. This hashing
24 value is mapped to the assigned DAM portion(s) 2202. The DAM portion 2202

1 that first added the atomic entry 2304 then replicates the atomic entry 2304 to the
2 assigned DAM portion(s) 2202.

3 With hashing for location selection, at least one DAM portion 2202 that has
4 a desired atomic entry 2304 locally cached at its DAM table 2206 is knowable
5 from the session identifier. A query atom 2204(Q) function call can therefore be
6 directed to the known DAM portion(s) 2202. This usually reduces network traffic
7 and/or latency.

8 This hashing for location selection may be used with one, two, three, or
9 more levels of replication with each range of hashing values mapping to one, two,
10 three, etc. different DAM portions 2202, respectively. Additionally, hashing for
11 location selection may be used with local caching upon acquisition.

12 FIG. 24 is a flow diagram 2400 that illustrates an exemplary method for
13 managing session information at multiple load balancing units. Flow diagram
14 2400 includes eight blocks 2402-2416. Although the actions of flow diagram
15 2400 may be performed in other environments and with a variety of software
16 schemes, FIGS. 1-3, 19, 20, 22, and 23B are used in particular to illustrate certain
17 aspects and examples of the method.

18 At block 2402, an incoming connection request with a session reference is
19 analyzed. For example, traffic routing functionality 2012 may receive an
20 incoming connection request that references a previously-opened/established
21 session of a particular type. At block 2404, a local DAM table is searched using
22 the session reference. For example, for a given load balancing unit 106 and traffic
23 routing functionality 2012, the DAM portion 2202 thereof may search its
24 corresponding DAM table 2206 looking for the session reference.

1 At block 2406, it is determined if the session reference matches a key of the
2 local DAM table. For example, DAM portion 2202 may search key fields
3 2304(K) of multiple entries 2304 of DAM table 2206 to determine whether the
4 session reference matches any values of the key fields 2304(K). If so, flow
5 diagram 2400 continues at block 2412.

6 If, on the other hand, the session reference does not match any key, flow
7 diagram 2400 continues at block 2408. At block 2408, a query atom function call
8 is made. For example, DAM portion 2202 may make a query atom 2204(Q)
9 function call that includes the session reference/identifier as the key. The query
10 atom 2204(Q) function call may be sent to at least one other DAM portion 2202.
11 The number, selection, order, etc. of possible destination DAM portions 2202 for
12 query atom 2204(Q) may depend on the options (e.g., replication level, hashing for
13 location selection, local caching upon acquisition, point-to-point versus multicast,
14 etc.) employed by DAM 2202.

15 At block 2410, a returned atom is received. For example, information from
16 a returned atom 2204(R) function call that is issued by another DAM portion 2202
17 may be received. The other DAM portion 2202 successfully located an atomic
18 entry 2304 in its corresponding DAM table 2206, with the located atomic entry
19 2304 having a key that matches the session reference. The information from the
20 returned atom 2204(R) function call includes values from key field 2304(K) and
21 data field 2304(D) for the located atomic entry 2304. These values correspond to
22 the session identifier of the session and the network address of the host 108 that is
23 affinitized to the session.

24 At block 2412, an atomic entry is extracted. The atomic entry is extracted
25 from the local DAM table if a match was found locally (at blocks 2404 and 2406)

1 or from the returned atom if a match was found elsewhere (at blocks 2408 and
2 2410). For example, an atomic entry 2304 may be extracted from DAM table
3 2206 of the DAM portion 2202 or from information received by a return atom
4 2204(R) function call. The extracted atomic entry 2304 may be cached at the local
5 DAM table 2206 if received as a result of the return atom 2204(R) function call.

6 At block 2414, the host having session affinity with the referenced session
7 is ascertained from the atomic entry. For example, a value of the data field
8 2304(D) of the extracted atomic entry 2304 may be ascertained to thereby
9 ascertain a network address of the affinitized host 108. At block 2416, the
10 incoming connection request is routed to the ascertained host. For example, traffic
11 routing functionality 2012 and/or forwarding functionality may route the incoming
12 connection request having the session reference to the ascertained and affinitized
13 host 108. Exemplary classifying, request routing, and forwarding functionalities
14 are described in the following section.

15 **Exemplary Classifying, Forwarding, and Request Routing**

16 This section describes how traffic routing may be implemented for network
17 load balancing, including with regard to high availability of such traffic routing
18 functionality. Traffic routing functionality may include classifying and/or
19 requesting routing functionality, especially in conjunction with forwarding
20 functionality. This section primarily references FIGS. 25-31. It illuminates the
21 functionality of a request router 306 (of FIG. 3), an interrelationship between
22 tracking sessions and utilizing health and load information when routing traffic,
23 operational implementations for traffic routing interactions with session
24 information and/or health and load information, failover procedures for high
25 availability of network load balancing infrastructure (including handling failures

1 of classifying, forwarding, and/or request routing components), additional network
2 load balancing infrastructure configurations, and so forth.

3 FIG. 25 illustrates exemplary network load balancing infrastructure having
4 request routing functionality as realized by request router 306(H/S). As noted
5 above with reference to traffic routing functionality 2012, traffic routing may
6 relate to classifying (e.g., with forwarding) and/or requesting routing. Packet-level
7 classifying, in conjunction with forwarding, is described above with particular
8 reference to FIG. 4. Request routing is described here with particular reference to
9 FIG. 25.

10 Request-level routing occurs at a higher level than that of packet-level
11 routing. Generally, a request router 306 acts as a proxy for an application 316
12 running on a host 108. Request router 306 terminates TCP connections, parses
13 (perhaps partially) each request from a client 102, and resubmits each request to
14 host 108. Request router 306 may perform pre-processing on the connection, such
15 as SSL decryption. Also, request router 306 may chose to absorb certain requests
16 (e.g., the request router may maintain a cache of responses), and it may
17 “arbitrarily” modify requests before forwarding them to hosts 108.

18 Request routers 306 are usually application-specific, and they may be rather
19 open-ended in what they are capable of doing. By way of example only, a single
20 class of request routers 306 – HTTP/SSL request routers 306(H/S) – are addressed
21 in the following description. As illustrated, a client 102 having a network address
22 C1 is communicating across network 104 with hosts 108(1) and 108(2) having
23 network addresses H1 and H2, respectively. The communications are effectuated
24 via load balancing infrastructure that includes an HTTP/SSL request router
25 306(H/S).

1 HTTP/SSL request router 306(H/S) terminates HTTP and SSL traffic,
2 decrypts SSL traffic, examines each HTTP request from client 102, applies
3 application-specific rules to classify each request and to determine the “best”
4 endpoint for that request while taking into account application endpoint health and
5 load information, and submits the request to the endpoint. The request submission
6 to the endpoint uses a separate TCP connection than that of the one originated by
7 client 102 (the latter connection is terminated at HTTP/SSL request router
8 306(H/S)). These actions may be considered as logically equivalent to the actions
9 performed by a classifier 304, but a difference arises in that these actions in
10 HTTP/SSL request router 306(H/S) are occurring at the logical request level for
11 each request within the TCP connection. HTTP/SSL request router 306(H/S), and
12 request routers 306 generally, can use the same (i) application health and load and
13 (ii) session tracking infrastructure that is used by classifiers 304.

14 HTTP/SSL request router 306(H/S) is acting as an intermediary between
15 client 102 and two hosts 108(1) and 108(2). It is handling two requests from client
16 102 over a single TCP connection. In a described implementation, the resulting
17 request routing involves a number of actions. First, client 102 establishes an http
18 or https connection [1] to HTTP/SSL request router 306(H/S) and sends a request
19 #1 2502(1).

20 Second, HTTP/SSL request router 306(H/S) terminates the SSL session (if
21 the traffic is SSL encrypted), parses request #1 2502(1), and examines the content
22 of request #1 2502(1). Taking into account application health and load as well as
23 session information, HTTP/SSL request router 306(H/S) determines that host
24 108(1) is the “best” host for this particular request #1 2502(1) in this example.
25

1 Third, HTTP/SSL request router 306(H/S) establishes a secondary TCP
2 connection [2] to host 108(1). This secondary TCP connection is not sourced from
3 a VIP address on network 104; instead, it is sourced from an address (not shown in
4 FIG. 25) that is dedicated to request router 306(H/S) to ensure that responses 2504
5 from host(s) 108 reach the correct request router 306. (There may be multiple
6 request routers 306 that are active even though one request router 306(H/S) is
7 shown in FIG. 25 for clarity.) It may alternatively use an existing connection [2] to
8 host 108(1). HTTP/SSL request router 306(H/S) then sends an e.g. unencrypted
9 version of request #1 2502(1) to host 108(1). Fourth, host 108(1) replies with a
10 response #1 2504(1). Fifth, HTTP/SSL request router 306(H/S) encrypts this
11 response #1 2504(1) and sends it back to client 102 on TCP connection [1].

12 Sixth, client 102 sends another request, request #2 2502(2). Request #2
13 2502(2) is handled similarly to the handling of request #1 2502(1), except that
14 HTTP/SSL request router 306(H/S) selects host 108(2). The different selection
15 may be because host 108(1) is now failing or more-heavily loaded, because
16 request #2 2502(2) is directed to a different URL than request #1 2502(1), and so
17 forth. Regardless, HTTP/SSL request router 306(H/S) establishes another
18 secondary TCP connection, but this secondary TCP connection [3] is to host
19 108(2). Unencrypted request #2 2502(2) is routed to host 108(2), and a response #2
20 2504(2) is received therefrom as a result. An encrypted version of response #2
21 2504(2) is then sent to client 102 from HTTP/SSL request router 306(H/S).

22 Seventh, client 102 closes TCP connection [1] with HTTP/SSL request
23 router 306(H/S). HTTP/SSL request router 306(H/S) (at some future time) closes
24 connections [2] and [3] that it made to hosts 108(1) and 108(2), respectively, on
25 behalf of client 102. TCP connection [2] may alternatively be closed after

1 HTTP/SSL request router 306(H/S) decides to open/use TCP connection [3] for
2 request #2 2502(2).

Because an HTTP/SSL request router 306(H/S) terminates the HTTP/HTTPS connection, HTTP/SSL request router 306(H/S) can do more than route requests. For example, HTTP/SSL request router 306(H/S) can potentially maintain its own cache of responses (e.g., with an out-of-band mechanism to invalidate the cache). As noted in the above example, HTTP/SSL request router 306(H/S) can also potentially route different kinds of requests to different sets of hosts 108 based on e.g. the requested URL. Conversely, HTTP/SSL request router 306(H/S) can potentially aggregate requests from many short-lived client connections and send them over a few, long-standing TCP connections to hosts 108. Such connection aggregation can reduce the TCP connection processing overhead in hosts 108.

Request routers of other classes may correspond to other exemplary protocols besides HTTP. For example, a request router may be a SOAP request router. SOAP request routers function analogously to an HTTP/SSL request router 306(H/S). However, SOAP request routers specialize in routing SOAP traffic. SOAP request routers understand SOAP headers and make routing decisions based on the SOAP headers as well as application health and load.

Both packet-level classification and forwarding (or packet-level routing) and request-level routing can provide some form of layer-7 load balancing. Layer-7 load balancing is described further below in the section entitled “Exemplary Connection Migrating with Optional Tunneling and/or Application-Level Load Balancing”. Packet-level routing provides read-only access to the initial portion of

1 a client's TCP connection data, and request-level routing provides read and modify
2 access to an entire data stream.

3 Packet-level routing typically has several advantages over request-level
4 routing. These advantages include transparency (client packets are delivered to
5 hosts as-is, preserving source and destination IP addresses and port numbers), low
6 processing overhead (generally, forwarding traffic involves a route lookup), low
7 latency (individual packets are forwarded, and packets are not queued once the
8 TCP connection destination has been determined), and high-availability (generally,
9 a failure in a forwarder does not terminate the TCP connection). Request-level
10 routing, on the other hand, typically has the following advantages over packet-
11 level routing: an ability to examine an entire data stream flowing to and from the
12 client; and an ability to transform a data stream, and even to split the data stream
13 among multiple hosts or aggregate data streams from multiple clients.

14 FIG. 26 is a flow diagram 2600 that illustrates an exemplary method for
15 routing incoming packets with regard to (i) session information and (ii) health and
16 load information. Flow diagram 2600 includes eight blocks 2602-2616. Although
17 the actions of flow diagram 2600 may be performed in other environments and
18 with a variety of software schemes, FIGS. 1-3, 12, 18-20, 22, and 23B are used in
19 particular to illustrate certain aspects and examples of the method.

20 At block 2602, an incoming packet is received. For example, a packet from
21 a client 102 may be received at a forwarder 302 of a load balancing unit 106. At
22 block 2604, it is determined if the received packet is for a preexisting session. For
23 example, forwarder 302 may consult a local DAM table 2206() to determine that
24 the received packet is already part of a TCP/IP session.

1 Additionally, forwarder 302 may consult the local DAM table 2206() and
2 determine that the received packet is not already part of a TCP/IP session. In this
3 case, forwarder 302 provides the received packet to a classifier 304, which checks
4 for a higher level session affinity for the received packet if it has a session
5 reference. Examples for these actions are described above with particular
6 reference to FIG. 24 and further below with particular reference to FIGS. 27 and
7 28.

8 If the received packet is for a preexisting session (as determined at block
9 2604), then flow continues at block 2606. At block 2606, a host that is affinitized
10 to the preexisting session is ascertained. For example, an affinitized host 108 may
11 be ascertained from the local DAM 2206() and/or the overall distributed DAM
12 2206 by forwarder 302 or classifier 304.

13 At block 2608, it is determined if the affinitized host is healthy. For
14 example, classifier 304 may consult a consolidated health and load cache 1208 to
15 determine if the affinitized host 108 is healthy, especially for those received
16 packets that are part of sessions that are of a higher logical level than TCP/IP
17 sessions. The action(s) of this block may be accomplished in conjunction with a
18 health and load handler 314.

19 If the affinitized host is healthy (as determined at block 2608), then flow
20 continues at block 2610. At block 2610, the received packet is routed to the
21 affinitized host. For example, forwarder 302 (for TCP/IP sessions) or classifier 304
22 (for higher-level sessions) may route the packet to the affinitized host 108. In an
23 alternative implementation, classifier 304 may return the received packet to
24 forwarder 302 for routing to the affinitized host 108 even for received packets that
25 are part of higher-level sessions.

If, on the other hand, the affinitized host is not healthy (as determined at block 2608), then flow continues at block 2612. Also, if on the other hand, the received packet is not for a preexisting session (as determined at block 2604), then flow continues at block 2612. At block 2612, a host is selected responsive to health and load information. For example, classifier 304 may select a host 108 from and/or using a health and load-related application allotment (e.g., from a target application endpoint allotment response 1804) that is attained from health and load handler 314. Examples for these action(s) are described above with particular reference to FIGS. 19 and 18 and further below with particular reference to FIG. 30.

At block 2614, the received packet is routed to the selected host. For example, classifier 304 may route (optionally via forwarder 302) the packet to the selected host 108. At block 2616, a route for a connection path to the selected host is plumbed. For example, classifier 304 may add a session information entry to DAM table 2206, especially at the DAM table 2206() that is local to the forwarder 302 that provided the received packet to the classifier 304. This session information entry may be replicated in accordance with the instituted redundancy policy for a DAM 2202 (e.g., of a session tracker 308).

The action(s) of block 2614 and those of block 2616 may be performed in the order specifically illustrated, with those of block 2616 being performed prior to those of block 2614, with the actions partially or fully overlapping in any order, and so forth. It should be noted that the actions performed by classifier 304 as described above may alternatively be performed by a request router 306 (or more generally traffic routing functionality 2012).

1 In addition to packet-level and request-level routing, traffic routing
2 functionality as described herein (e.g., traffic routing functionality 2012, a request
3 router 306, a forwarder 302/classifier 304 pair, etc.) can also be used to implement
4 firewall functionality. Hence, a feature of the traffic routing functionality may
5 include blocking traffic, instead of automatically routing traffic to the correct host
6 108. For example, a classifier 304 can inspect traffic and drop it if it is deemed
7 unsafe.

8 FIG. 27 illustrates an exemplary traffic routing flow in the absence of
9 failures. As illustrated, one or more load-balancing-aware switches 202(LBA)
10 front the remaining load balancing infrastructure 106 (not separately indicated).
11 Forwarding and classifying functionality are distributed across three devices or
12 nodes. A first device includes forwarder 302(1) and classifier 304(1). A second
13 device includes classifier 304(2). A third device includes forwarder 302(2).

14 With classifier 304(2) executing on the second device and forwarder 302(2)
15 executing on the third device, each device may be specially tuned for its respective
16 functions. For example, the hardware, software, firmware, some combination
17 thereof, etc. of the second device and the third device may be adapted to support
18 the desired functionality without excessive over provisioning. Thus, the third
19 device that includes forwarder 302(2) may be akin to a switch and/or router from a
20 hardware capability perspective, and the second device that includes classifier
21 304(2) may be more akin to a server and/or personal computer from a hardware
22 capability perspective.

23 Although shown as three devices that are providing functionality across
24 four components, alternative logical and/or device-level configurations for
25 forwarding and classifying functionality are applicable to the exemplary traffic

1 routing flow that is described here for FIG. 27. Also, although the routing
2 destinations are shown as hosts 108, the descriptions herein of routing
3 implementations may alternatively be applied more generally to a next node
4 destination for the packet and not necessarily a final node that consumes the
5 packet.

6 A DAM 2202 realization of session tracker 308 is used to implement DAM
7 table 2206. However, session affinity preservers 1904 in general are also
8 applicable to the exemplary traffic routing flow of FIG. 27. Forwarder 302(1)
9 includes DAM table portion 2206(1), and forwarder 302(2) includes DAM table
10 portion 2206(2). Incoming packets are routed to host 108(1) or host 108(2).

11 In a described implementation, DAM 2202 is a distributed, in-memory
12 table of “atoms” 2304 (e.g., keyword-value pairs, with optional metadata) having
13 session information. DAM 2202 and DAM table 2206 is described further above
14 with particular reference to FIGS. 22-24. Any node in the cluster of classifiers 304
15 may add, query, and delete atoms 2304. DAM 2202 maintains a highly available
16 DAM table 2206 that includes active (e.g., TCP/IP level) routes as well as higher-
17 level session information. Examples of higher level sessions include: a TLS/SSL
18 session, a PPTP session, an IPSec/L2TP session, an ISA session, an HTTP cookie-
19 based session, and so forth. Furthermore, DAM 2202 may include session
20 information entries in DAM table 2206 that are directed to other non-TCP/IP
21 sessions, such as RTP, UDP, and so forth.

22 At (1), load-balancing-aware switches 202(LBA) direct an incoming packet
23 to forwarder 302(1). At (2), forwarder 302(1) consults its internal routing table,
24 DAM table 2206(1). When forwarder 302(1) does not find an atomic entry 2304
25

1 for this packet, it forwards the packet to its assigned and/or associated classifier,
2 classifier 304(1).

3 At (3), classifier 304(1) recognizes that the packet in this example is a first
4 packet of a new session (e.g., a SYN packet for a TCP connection). Classifier
5 304(1) therefore treats the packet as a start of a new TCP connection from a client
6 102. Using health and load information from a health and load handler 314 (not
7 explicitly illustrated), classifier 304(1) determines that host 108(1) should receive
8 this session.

9 Classifier 304(1) updates DAM table 2206(1) that serves as the local
10 routing table for forwarder 302(1), and it also inserts an atomic entry 2304
11 representing the route into the overall DAM 2206. These may be separate
12 operations, a single operation in which the TCP/IP-level sessions of DAM table
13 2206 are located at forwarders 302, and so forth. DAM 2202 internally replicates
14 this route to one or more other members of the cluster of classifiers 304 in
15 accordance with its stipulated redundancy policy. Classifier 304(1) may
16 optionally communicate with host 108(1) to confirm the creation of the new
17 session before it updates DAM table 2206(1) of forwarder 302(1) and the overall
18 DAM 2202/DAM table 2206.

19 At (4), forwarder 302(1) directly forwards subsequent packets for this
20 connection to host 108(1) without interacting with classifier 304(1). DAM 2202
21 can be used to mask, at least in part, the failure of a forwarder 302, a classifier
22 304, or a forwarder/classifier pair 302/304. DAM 2202 can also be used, at least
23 in part, to preserve client connectivity if load-balancing-aware switches 202(LBA)
24 inadvertently start sending packets for an established connection to a different
25 forwarder 302.

FIG. 28 illustrates an exemplary traffic routing flow in the presence of failure(s). In contrast to the “failure-free” exemplary traffic routing flow of FIG. 27, a failure has occurred in a portion of the network load balancing infrastructure 106 (not specifically identified) of FIG. 28. Specifically, the first device, on which forwarder 302(1) and classifier 304(1) are resident and executing, fails after the connection that is illustrated in FIG. 27 is established. This failure is masked, at least in part, by DAM 2202.

At (1), load-balancing-aware switches 202(LBA) detect the failure of forwarder 302(1) and start forwarding packets for the connection to some other forwarder 302 in the cluster. In this example, the other forwarder 302 is forwarder 302(2). Although FIG. 28 illustrates a failure situation, load-balancing-aware switches 202(LBA) may also send this traffic to forwarder 302(2) even if forwarder 302(1) is still available. This non-failure-induced change of forwarders 302 may occur, for example, because load-balancing-aware switches 202(LBA) do not preserve the affinity of this traffic to forwarder 302(1). Any of several factors can cause switches 202 to (mis)direct traffic to a different, non-affinitized forwarder 302. For example, traffic for the same higher-level session can arrive at switches 202 from a different source IP address or source port when the source is behind a farm of proxy servers. The actions of notations (2)-(5) apply to both the failure and the “misdirected traffic” situations.

At (2), forwarder 302(2) consults its routing table, DAM table 2206(2). When it does not find a route for this packet, it forwards the packet to its classifier 304(2). At (3), classifier 304(2) recognizes that this packet is a “mid-session” packet, and classifier 304(2) queries DAM 2202 for the route for this packet.

1 DAM 2202 responds with the route for the connection from an atomic entry 2304
2 that is associated therewith.

3 At (4), classifier 304(2) plumbs the route in forwarder 302(2). An
4 exemplary protocol for plumbing routes is described further below. At (5),
5 subsequent packets for this connection that are directed to forwarder 302(2) are
6 routed directly to the correct host, which is host 108(1) in this example, without
7 consulting classifier 304(2).

8 Generally, a route plumbing protocol for communications between
9 classifiers 304 and forwarders 302 includes instructions to add and remove routes.
10 More specifically, an add route instruction is sent from a classifier 304 to a
11 forwarder 302 in order to plumb a route from the forwarder 302 to a destination
12 host 108 for a given connection. By way of example, an add route instruction can
13 be provided to forwarder 302(2) from classifier 304(2) as indicated at (4) in FIG.
14 28. The route (e.g., a key and corresponding value) is added to local DAM table
15 2206(2) for quick access by forwarder 302(2) in the future. In this example,
16 classifier 304(2) is a separate device from forwarder 302(2), so the route plumbing
17 protocol may be an inter-device protocol. However, the route plumbing protocol
18 may also be utilized for intra-device communications.

19 In a described implementation, classifier 304(2) includes a connection
20 inventory 2802. With connection inventory 2802, classifier 304(2) keeps track of
21 the sessions of any forwarders 302 (such as forwarder 302(2)) for which classifier
22 304(2) plumbs routes. To enable classifier 304(2) to keep track of the sessions,
23 including cessations thereof, forwarder 302(2) forwards final packets for sessions
24 (such as a TCP FIN packet) to classifier 304(2). Classifier 304(2) then deletes an
25 entry in connection inventory 2802 that corresponds to the session and sends a

1 delete route instruction to forwarder 302(2). Upon receiving the delete route
2 instruction, forwarder 302(2) removes the corresponding route in DAM table
3 2206(2).

4 In this manner, the classifying functionality in conjunction with session
5 tracking functionality can control the route tables, and the routes thereof, that are
6 used by the forwarding functionality. Consequently, forwarding functionality that
7 is separated onto a different device may be effectuated using high-speed, but
8 relatively simple, hardware. Alternatively, classifiers 304 may rely on
9 communications with/from hosts 108, rather than (or in addition to) intercepted
10 session initiation (such as TCP SYN) and termination (such as TCP FIN) packets,
11 to determine the lifetimes of sessions. In other words, classifiers 304 may
12 alternatively or additionally receive and utilize session (up/down) messages
13 2008(U/D) as described above in the section entitled “Exemplary Session
14 Tracking”.

15 FIG. 29 illustrates additional exemplary failover procedures for high
16 availability of network load balancing infrastructure 106. Failover procedures for
17 two different failures, failure 2902 and failure 2906, are described. As illustrated,
18 network load balancing infrastructure 106 (not separately indicated) includes five
19 components: forwarder 302(1), forwarder 302(2), forwarder 302(3), classifier
20 304(1), and classifier 304(2).

21 In a described implementation, each of these five components 302(1),
22 302(2), 302(3), 304(1), and 304(2) corresponds to an individual device. However,
23 similar failover procedures apply to environments in which different load
24 balancing components share devices. Also, similar or analogous failover

25

1 procedures may apply to environments having other numbers, combinations,
2 scalings, etc. of components.

3 Initially at [1], router/switch(es) 202 direct an incoming packet that
4 happens to be for a new connection to forwarder 302(1). Because forwarder
5 302(1) does not have a route for this connection in its local routing table, it sends
6 the packet to classifier 304(1) as indicated by the dashed double arrow at (1).
7 Classifier 304(1) first checks session information with reference to session
8 tracking 308 for a possible higher-level session affinity. In this example, the
9 packet is not affinized to an existing session, so classifier 304(1) selects a host 108
10 with reference to health and load information with reference to health and load
11 handling 314.

12 Specifically, classifier 304(1) selects host 108(1) in this example.
13 Assuming the packet is for a TCP/IP connection, this TCP/IP session as linked to
14 host 108(1) is added to DAM 2202 using an add atom 2204(A) function call by
15 classifier 304(1). The initial packet is forwarded to host 108(1) by classifier
16 304(1) or forwarder 302(1). Classifier 304(1) also plumbs a route in the local
17 routing table of forwarder 302(1). Subsequent packets are forwarded to host
18 108(1) by forwarder 302(1) without further interaction with classifier 304(1).

19 At some time during connection [1], there is a failure 2902 at forwarder
20 302(1). With load-balancing-aware router/switch(es) 202(LBA), this failure 2902
21 is detected. As a result, at point 2904, router/switch(es) 202 direct later packets
22 that would have been sent to forwarder 302(1) along connection [1] to another
23 forwarder 302, which is forwarder 302(2) in this example.

24 Forwarder 302(2) thus receives future packets along a connection [2].
25 Because forwarder 302(2) does not have an entry in its local routing table for the

1 packets that were formerly directed to forwarder 302(1), forwarder 302(2) sends
2 the first received packet of connection [2] to the classifier to which it is
3 assigned/associated. In this example, forwarder 302(2) is assigned to classifier
4 304(2) as indicated by the dashed double arrow at (2).

5 Classifier 304(2) uses a query atom 2204(Q) function call to attain the
6 atomic entry 2304 (not explicitly shown) from DAM 2202 that is associated with
7 the existing TCP/IP connection. This atomic entry 2304 is provided through DAM
8 2202 of session tracking 308 via a return atom 2204(R) function call. Classifier
9 304(2) extracts the host 108(1) that is affinitized with this TCP/IP connection from
10 the returned atomic entry 2304. Classifier 304(2) forwards the first received
11 packet for connection [2] to host 108(1) and also plumbs a route in the local
12 routing table of forwarder 302(2). Subsequent packets are forwarded to host
13 108(1) by forwarder 302(2) without further interaction with classifier 304(2).

14 The above descriptions focus predominantly on failures of individual
15 forwarder 302 components. However, classifier 304 components can also fail.
16 For example, at some point, there is a failure 2906 at classifier 304(2). Forwarder
17 302(2) detects failure 2906 when it attempts to consume classification services or
18 through noticing a lack of some aliveness indication such as a heartbeat-type
19 indicator. To handle failure 2906, forwarder 302(2) is reassigned or re-associated
20 with a different classifier 304, which is classifier 304(1) in this example. Future
21 classification functionality is provided to forwarder 302(2) by classifier 304(1) as
22 indicated by the dashed double arrow at (3).

23 FIG. 30 illustrates an exemplary operational implementation of traffic
24 routing interaction with health and load information. Forwarder 302 and classifier
25 304 interact with health and load handler 314 in order to route packets to hosts

1 108(1), 108(2) ... 108(n). Although a forwarder 302 and a classifier 304 are
2 illustrated, the exemplary operational implementation is also applicable to a
3 request router 306 (or traffic routing functionality 2012 in general).

4 As illustrated, host 108(1) includes application endpoints IP1, IP3, and IP4
5 for application #1, application #1, and application #2, respectively. Host 108(2)
6 includes application endpoints IP2 and IP6 for application #1 and application #2,
7 respectively. Host 108(n) includes application endpoint IP5 for application #2.
8 These hosts 108(1), 108(2) ... 108(n) and application endpoints IP1, IP2, IP3, IP4,
9 IP5, and IP6 are monitored by health and load handler 314 (e.g., using health and
10 load infrastructure 1202, consolidated health and load cache 1208, etc.).

11 In a described implementation, at (1) classifier 304 requests one or more
12 application endpoint allotments (e.g., via at least one target application endpoint
13 allotment request 1802) in an environment using a token allotment scheme 1806.
14 Health and load handler 314, in this example, responds by providing token
15 allotments 3002 (e.g., via at least one target application endpoint allotment
16 response 1804).

17 Specifically, a token allotment for application #1 3002(1) and a token
18 allotment for application #2 3002(2) are available to classifier 304. Token
19 allotment for application #1 3002(1) initially provides 40 tokens for IP1, 35 tokens
20 for IP2, and 25 tokens for IP3. Token allotment for application #2 3002(2)
21 provides 10 tokens for IP4, 72 tokens for IP5, and 18 tokens for IP6. For each
22 new connection that is allocated a routing to an application endpoint by classifier
23 304, a token is consumed by classifier 304.

24 At (2), forwarder 302 receives an initial incoming packet for a new
25 connection. Because no routing for this new connection is present in local DAM

1 table portion 2206 of forwarder 302, forwarder 302 forwards the initial packet to
2 classifier 304 at (3).

3 At (4), classifier 304 (e.g., after determining that the initial packet does not
4 include a session reference for a higher-level session) selects an application
5 endpoint (and thus a host 108) responsive to health and load information.
6 Specifically, for a new connection that is to be served by application #1, classifier
7 304 can select any of IP1, IP2, and IP3 if a token for the respective endpoint still
8 exists.

9 Classifier 304 can consume tokens in any of many possible manners. For
10 example, classifier 304 may use a round-robin approach regardless of the number
11 of tokens per endpoint. Alternatively, classifier 304 may simply start from IP1 and
12 progress through IP3 while consuming all tokens for each endpoint before moving
13 to the next endpoint in a linear approach. Also, classifier 304 may consume a
14 token from the endpoint-defined-set of tokens that currently has the greatest
15 number of tokens at any one moment. Using the latter approach, classifier 304
16 selects IP1. Other approaches may also be employed.

17 As illustrated, classifier 304 consumes a token for application endpoint IP2.
18 Consequently, the token set for IP2 is reduced from 35 tokens to 34 tokens as a
19 token is consumed. Also, the initial packet for the new connection is to be routed
20 to application endpoint IP2.

21 At (5A), the initial packet is forwarded from classifier 304 to application
22 endpoint IP2 of host 108(2). Before, during, or after this forwarding, classifier
23 304 at (5B) plumbs a route for this connection in local DAM table portion 2206.
24 Classifier 304 may also add an atomic entry 304 for this session into DAM table
25 2206 for distribution and replication purposes. At (6), future packets for this

1 connection/session are forwarded from forwarder 302 to application endpoint IP2
2 of host 108(2) using the local routing table of forwarder 302 as realized by local
3 DAM table portion 2206 in FIG. 30.

4 FIG. 31 illustrates exemplary high availability mechanisms for network
5 load balancing infrastructure 106. Specifically, exemplary failure detection 3104,
6 exemplary failure handling 3106, and exemplary failure recovery 3108 are shown.
7 These exemplary high availability mechanisms are described with regard to
8 different network load balancing infrastructure 106 components. The network
9 load balancing infrastructure 106 components include a forwarder 302, a classifier
10 304, a request router 306, a session tracker 308, and a health and load handler 314.

11 At 3102(A), forwarder 302 undergoes a local failure. At 3104(A), at least
12 one load-balancing-aware switch detects the failure. To handle local failure
13 3102(A), packets are redirected to other forwarder(s) at 3106(A) by the load-
14 balancing-aware switch. To recover from the failure of forwarder 302, routes that
15 were stored locally at forwarder 302 are rebuilt at 3108(A) at the forwarder(s) to
16 which packets are redirected using a distributed session tracking manager and a
17 table thereof such as a DAM and a DAM table thereof. The distributed session
18 tracking manager may therefore include data redundancies of one or more levels.

19 At 3102(B), classifier 304 undergoes a local failure. At 3104(B), at least
20 one forwarder detects the failure. To handle local failure 3102(B), packets are
21 redirected to other classifier(s) at 3106(B) by the forwarder detecting the failure.
22 To recover from the failure of classifier 304, session information that was stored
23 locally at classifier 304 are rebuilt at 3108(B) at the classifier(s) to which packets
24 are redirected using DAM. This session information may be, for example, session
25 information of a higher level than baseline TCP/IP connections. Also, such

1 session information may be considered as part of session tracking infrastructure
2 that is resident on the same device as classifier 304.

3 At 3102(C), request router 306 undergoes a local failure. At 3104(C), at
4 least one forwarder and/or load-balancing-aware switch detect the failure. To
5 handle local failure 3102(C), packets are redirected to other request router(s) at
6 3106(C) by the forwarder and/or load-balancing-aware switch. Individual current
7 logical requests on which request router 306 is working upon the occurrence of
8 local failure 3102(C) may be lost unless each such individual logical request is
9 replicated while the request is being serviced. To recover from the failure of
10 request router 306, session information and/or routes that were stored locally at
11 request router 306 are rebuilt at 3108(C) at the request router(s) to which packets
12 (and thus new logical requests) are redirected. The session information rebuilding
13 may be effectuated using DAM. Again, such session information may be
14 considered as part of session tracking infrastructure that is resident on the same
15 device as request router 306.

16 At 3102(D), session tracker 308 undergoes a local failure. At 3104(D), at
17 least one forwarder and/or classifier detect the failure. For example, if session
18 tracker 308 is resident on a same device as a classifier, then a forwarder or another
19 classifier may detect the failure. If session tracker 308 is resident on a separate
20 device, then a classifier may detect the failure. To handle local failure 3102(D),
21 data redundancy of one or more levels and distribution across multiple devices are
22 instituted at 3106(D) for the tracked session information. It should be noted that
23 the redundancy and distribution are instituted prior to failure 3102(D). To recover
24 from the failure of session tracker 308, session information from the tables of the
25 DAM may be redistributed and re-replicated at 3108(D) across at least two devices

1 (if not already so distributed and sufficiently replicated) in order to handle a
2 second level of failure.

3 At 3102(E), health and load handler 314 undergoes a local failure. At
4 3104(E), at least one classifier and/or request router detect the failure. For
5 example, a component that is receiving health and load information from health
6 and load handler 314 may detect a failure if health and load handler 314 becomes
7 non-responsive, especially if health and load handler 314 is resident on a different
8 device from that of the inquiring component. To handle local failure 3102(E),
9 cached health and load data redundancy and intrinsic failure handling are
10 employed at 3106(E) for the health and load information.

11 For example, each health and load handler 314 can include a consolidated
12 health and load information cache 1208 that duplicates information in health and
13 load tables 1204 on multiple hosts 108. Also, consumers of the health and load
14 information 1206 of a given health and load handler 314 may be located on a same
15 device as health and load handler 314 so that failure of health and load handler
16 314 is intrinsically acceptable. Similarly, the authoritative version of a respective
17 portion of health and load information 1206 is located on a respective host 108 so
18 that failure of the host 108 renders the loss of the respective portion of the health
19 and load information acceptable.

20 To recover from the failure of health and load handler 314, a given network
21 load balancing component that consumes health and load information may query a
22 different health and load handler because each such health and load handler
23 includes a consolidated cache of health and load handler information. Also, when
24 health and load handler 314 is again accessible, message protocol 1500 may be
25 used at 3108(E) to rebuild its consolidated cache of health and load information.

1 Using these exemplary high availability mechanisms, failures of network load
2 balancing infrastructure 106 components can be detected, handled, and recovered
3 from in order to mask such failures from clients 102.

4 **Exemplary Connection Migrating**

5 **with Optional Tunneling and/or Application-Level Load Balancing**

6 This section describes how connection manipulation, such as connection
7 migration, may be utilized in network load balancing. This section primarily
8 references FIGS. 32-39 and illuminates connection migrating functionality such as
9 that provided by connection migrator 310 (of FIG. 3). As described above with
10 reference to FIGS. 3 and 4, each incoming connection at load balancing
11 infrastructure 106 may be terminated thereat. Afterwards, the connection may be
12 migrated to a host 108 such that the connection is then terminated at the host 108.
13 Connection migrator 310 is capable of performing this connection migration and
14 may be located partially at hosts 108 to effectuate the migration. Such connection
15 migration may be performed in conjunction with application-level load balancing
16 by a classifier 304 and/or using tunneling via tunneler 312.

17 FIG. 32 illustrates an exemplary approach to application-level network load
18 balancing with connection migration. Application-level, or layer-7, load balancing
19 pertains to making load balancing decisions with regard to an application that is to
20 handle a connection. To perform application-level load balancing, load balancing
21 infrastructure 106 usually takes into consideration a data portion of a connection.
22 Unless request routing is employed, a classifier 304 typically takes a peek at the
23 initial portion of a connection and then migrates the connection, in conjunction
24 with connection migrator 310, to a selected host 108.

25

1 For application-level load balancing in a TCP-based environment generally,
2 classifiers 304 peek at the initial portion of a client's TCP data when deciding
3 where to forward the client's TCP connection. Thus, application-level logic
4 examines the client's data and makes load balancing decisions based on that data.
5 For example, if a connection is an (unencrypted) HTTP connection, a classifier
6 304 can take a peek at the HTTP header of the first HTTP request in the
7 connection, and it can make routing decisions based on some portion of the
8 content of the header (e.g., the URL, a cookie, etc.). Although application-level
9 load balancing, connection migration, and tunneling are applicable to other
10 protocols, TCP/IP is used predominantly in the examples herein.

11 As illustrated, load balancing infrastructure 106 (not specifically indicated)
12 includes a forwarder 302, a classifier 304, a tunneler 312, and a connection
13 migrator 310 (and possibly e.g. load-balancing-aware router/switches 202(LBA)).
14 Forwarder 302 corresponds to the virtual IP address and forwards packets to hosts
15 108 in accordance with host selections by classifier 304. Although not specifically
16 shown in FIG. 32 for clarity, hosts 108 also include connection migrator 310
17 functionality and tunneler 312 functionality.

18 In a described implementation, forwarder 302, classifier 304, and
19 connection migrator 310 (at classifier 304 and on hosts 108), along with TCP
20 protocol software on classifier 304 and hosts 108, cooperate to provide connection
21 migration. The connection migration illustrated in FIG. 32 is for a connection
22 from client 102(1) that is initially terminated at classifier 304. After connection
23 migration, the connection from client 102(1) is terminated at host 108(1). Once
24 the connection is terminated at host 108(1), packets for the connection may be
25 tunneled using tunneler 312 (at forwarder 302 and host 108(1)).

1 At (1), client 102(1) sends a SYN packet to forwarder 302 to signal the start
2 of a new TCP connection. At (2), forwarder 302 forwards this packet to classifier
3 304. At (3), classifier 304 accepts the TCP connection on behalf of a host 108
4 (whose identity is not yet known because the actual target host 108() has yet to be
5 selected). In TCP protocol terms, classifier 304 sends a SYN-ACK packet to
6 client 102(1).

7 At (4), client 102(1) begins sending data. (The initial SYN packet may also
8 contain data.) The data is processed by classifier 304, which can consult
9 application-specific logic. The application-specific logic can relate to which host
10 108 is capable of handling or best handling which types of requests or
11 connections. Hence, classifier 304 uses the data, as well as application health and
12 load information from health and load handler 314 and optionally application
13 session information from session tracker 308, to determine a host 108 that is better
14 or best suited to handle this connection from client 102(1). In this example, host
15 108(1) is selected.

16 At (5), classifier 304 sends a “binary blob” that represents the state of the
17 TCP connection to host 108(1). This connection state is aggregated with
18 cooperation from a TCP stack on classifier 304 by connection migrator 310. The
19 binary blob contains data from client 102(1) that has been acknowledged by
20 classifier 304 and TCP parameters such as the TCP/IP 4-tuple, initial sequence
21 numbers, and so forth.

22 At (6), a connection migrator 310 component on host 108(1) (not explicitly
23 shown in FIG. 32) “injects” this connection into a TCP stack on host 108(1) using
24 the state of the TCP connection from the binary blob received from classifier 304.
25 This connection state injection is performed in cooperation with the TCP stack on

1 host 108(1), making it appear to applications 316 on host 108(1) that this
2 connection was originally accepted by host 108(1) itself. Client 102(1) and
3 applications 316 on host 108(1) are unaware of the connection migration.

4 At (7), classifier 304, in cooperation with the TCP stack on classifier 304,
5 cleans up the internal state maintained for this connection. This internal state
6 cleanup at classifier 304 is performed silently such that client 102(1) is not
7 notified that the connection state is being torn down. Classifier 304 also adds a
8 route in a local routing table of forwarder 302 that indicates host 108(1) as the
9 destination for packets of this connection.

10 At (8), subsequent packets for the connection are routed by forwarder 302
11 to host 108(1) without diversion to or through classifier 304. These packets may
12 be treated the same by forwarder 302 as those packets for connections that are
13 classified and routed without using connection migration. These subsequent
14 packets may optionally be tunneled from forwarder 302 to host 108(1) using
15 tunneler 312. Tunneler 312 is also illustrated (using dashed lines) at connection
16 migrator 310 at classifier 304 because certain parameter(s) used by tunneler 312
17 may be determined during a connection migration and/or associated with a
18 connection being migrated. Exemplary implementations for tunneler 312 are
19 described further below with particular reference to FIGS. 38 and 39.

20 FIG. 33 is a flow diagram 3300 that illustrates an exemplary method for
21 migrating a connection from a first device to a second device. Flow diagram 3300
22 includes seven blocks 3302-3314. Although FIGS. 32 and 34-37 focus primarily
23 on connection migration in a network load balancing environment, connection
24 migration as described herein may be effectuated between two devices in general

1 that each include connection migration functionality, such as that of connection
2 migrator 310.

3 At block 3302, a connection is accepted at a first device. For example, a
4 first device may terminate an incoming connection in accordance with one or more
5 protocols of a protocol stack portion of a network stack. At block 3304, data is
6 received for the connection at the first device. For example, this data may be
7 received in an initial packet that requests the connection or in one or more packets
8 that are received subsequent to an acceptance of the connection.

9 At block 3306, a connection state for the accepted connection is aggregated
10 from a protocol stack (or more generally from a network stack) at the first device.
11 For example, a protocol state of the one or more protocols of the protocol stack
12 may be compiled and aggregated with any received data that has been
13 acknowledged. At block 3308, the connection state is sent from the first device to
14 a second device. For example, the aggregated information of the connection state
15 may be sent using a reliable protocol to a second device.

16 At block 3310, the connection state for the connection being migrated is
17 received from the first device at the second device. At block 3312, the connection
18 state is injected into a protocol stack (or more generally into the network stack) of
19 the second device. For example, the connection may be rehydrated using the
20 protocols of the protocol stack of the second device such that programs above the
21 protocol stack level are unaware that the connection is a migrated connection.
22 More specifically, the protocol state may be infused into the protocol stack. The
23 aggregated data of the connection state is also incorporated at the second device.
24 At block 3314, the connection is continued at the second device. For example, the
25

1 connection may be continued at the second device as if the connection was not
2 previously terminated elsewhere.

3 FIG. 34 illustrates an exemplary approach to connection migration from the
4 perspective of an originating device 3400. Connection migration in originating
5 device 3400 is effectuated, at least partly, by connection migrator 310. In a
6 described implementation, originating device 3400 is a device that is part of
7 network load balancing infrastructure 106. For example, originating device 3400
8 may comprise a classifier 304, possibly along with a forwarder 302, a request
9 router 306, and so forth.

10 As illustrated, originating device 3400 includes as parts of its network stack
11 a physical network interface (PNI) 3410, a PNI miniport 3408, a protocol-
12 hardware interface 3406, a protocol stack 3404, and a socket layer 3402.
13 Originating device 3400 also includes load balancing functionality 106, such as a
14 classifier 304 at an application level and connection migrator 310. Specifically,
15 connection migrator 310 includes a migrator intermediate driver 3414 and a
16 migrator shim 3412. Connection migrator 310 is capable of offloading a
17 connection from originating device 3400.

18 In a described implementation, physical network interface 3410 may be a
19 network interface card (NIC) (e.g., an Ethernet NIC), a wireless interface, and so
20 forth. Although only one physical network interface 3410 is shown, a given
21 device may actually have multiple such physical network interfaces 3410 (i.e.,
22 originating device 3400 may be multi-homed). Each physical network interface
23 3410 typically corresponds to one or more physical network addresses.

24 PNI miniport 3408 is a software module that understands and interfaces
25 with the specific hardware realization of physical network interface 3410.

1 Protocol-hardware interface 3406 is a layer that includes one or more respective
2 interfaces between one or more respective protocols and PNI miniport 3408.

3 Protocol stack 3404 includes one or more respective modules that are each
4 directed to one or more respective protocols. Examples of such protocols are
5 described further below with reference to FIGS. 36 and 37. In a transient context,
6 protocol stack 3404 includes a protocol state 3420 for each connection existing at
7 originating device 3400. A socket layer 3402 lies between a program such as load
8 balancing functionality 106 and protocol stack 3404. Socket layer 3402 provides
9 APIs between load balancing functionality 106 and protocol stack 3404, and it
10 enables programs to register for connections, among other things.

11 Migrator intermediate driver 3414, or more generally migrator driver 3414,
12 is located at protocol-hardware interface layer 3406. Migrator shim 3412 is
13 located transparently between protocol stack 3404 and socket layer 3402.

14 When an initial packet (not shown) requesting a new connection is
15 presented to originating device 3400, the packet is directed upward from physical
16 network interface 3410, to PNI miniport 3408, through protocol-hardware
17 interface layer 3406, and to protocol stack 3404. As the packet traverses the one
18 or more protocols of protocol stack 3404, protocol state 3420 is created thereat.
19 Also, as a result of this initial packet or as a consequence of load balancing
20 functionality 106 accepting the connection to take a peek at the request, data 3416
21 arrives at originating device 3400.

22 In operation, migrator intermediate driver 3414 diverts a copy of data 3416
23 to the logic of connection migrator 310. When load balancing functionality 106
24 issues a migrate connection function call, the migrate function call is passed to a
25 topmost layer of protocol stack 3404 so that connection state aggregation 3418

1 may commence. Protocol state 3420 is compiled from the one or more protocols
2 of protocol stack 3404. In a TCP/IP implementation, protocol state 3420 may
3 include (i) destination and source TCP ports and IP addresses (e.g., a TCP/IP 4-
4 tuple), (ii) TCP window state, (iii) initial sequence numbers, (iv) timeout
5 information, (v) IP fragment ID, (vi) routing information, and (vii) so forth.

6 Connection state aggregation 3418 also aggregates data 3416 that has been
7 diverted to connection migrator 310 and that has already been acknowledged from
8 originating device 3400 (e.g., by load balancing functionality 106). This
9 aggregated connection state 3418 includes protocol state 3420 and data 3416 (and
10 optionally other connection-related information). Aggregated connection state
11 3418 is then sent as a binary blob 3422 away from originating device 3400 toward
12 a targeted device.

13 Binary blob 3422 may be sent from originating device 3400 toward a
14 targeted device using a reliable protocol. “Reliable” may imply, for example, that
15 binary blob 3422 is received intact at the targeted device even if individual packets
16 that constitute binary blob 3422 are lost or corrupted. This binary blob 3422 may
17 also be bundled with a flow identifier if the connection is to be tunneled
18 subsequently with tunneler 312. Flow identifiers with tunneling are described
19 further below with particular reference to FIGS. 38 and 39.

20 FIG. 35 illustrates an exemplary approach to connection migration from the
21 perspective of a target device 3500. Target device 3500 is similar to originating
22 device 3400 with respect to the various illustrated layers/modules, including
23 connection migrator 310. As illustrated however, at least one application 316 at an
24 application level is interfacing with socket layer 3402. Target device 3500 may

1 therefore comprise a host 108. Also, connection migrator 310 is capable of
2 uploading a connection from originating device 3400.

3 In a described implementation, application 316 is the destination of the
4 connection-initiating packet received at originating device 3400. From originating
5 device 3400, target device 3500 receives binary blob 3422. Binary blob 3422
6 includes the connection state associated with the connection being migrated to
7 target device 3500 and optionally a flow identifier. This connection state includes
8 protocol state 3420 and acknowledged data 3416 (and possibly other connection-
9 related information).

10 In operation, when binary blob 3422 reaches protocol-hardware interface
11 layer 3406, migrator intermediate driver 3414 recognizes it as a blob for
12 connection migration and diverts it. The connection state is injected at 3502 to
13 create the appearance to application 316 that the connection was originally
14 terminated at target device 3500.

15 Specifically, protocol state 3420 of injected connection state 3502 is infused
16 into protocol stack 3404. In a described implementation, protocol state 3420 is
17 infused first at higher-level protocols and then at lower-level protocols of protocol
18 stack 3404. After protocol state 3420 is infused into protocol stack 3404, data
19 3416 can be indicated up to application 316. This data 3416 can be provided to
20 application 316 as if it were part of a newly and locally terminated connection.

21 After connection state injection 3502 is completed, the connection initiated
22 by the packet received at originating device 3400 is successfully migrated
23 therefrom to target device 3500. Subsequent packets for the connection may be
24 forwarded directly to target device 3500 without passing through originating
25 device 3400, or at least with only simple routing and no application-level analysis

1 being applied thereto. Optionally, these packets may be tunneled such that
2 migrator intermediate driver 3414 effectively operates as a software-based virtual
3 NIC that is bound to the virtual IP address. In other words, migrator intermediate
4 driver 3414 (of FIG. 35) may comprise a virtual network adapter that is bound to
5 the destination address of un-encapsulated packets.

6 FIG. 36 illustrates an exemplary approach to an offloading procedure 3600
7 for a connection migration. Migration offloading procedure 3600 illustrates
8 additional exemplary details for a connection migration by an originating device
9 3400. As illustrated, general protocol stack 3404 includes a TCP stack 3404(T), an
10 IP stack 3404(I), and an address resolution protocol (ARP) stack 3404(A).
11 However, other specific protocol stacks 3404() may alternatively be employed.

12 By way of example, protocol-hardware interface layer 3406 may be
13 realized as a network driver interface specification (NDIS)-based layer in a
14 Microsoft® Windows® operating system (OS) environment. Also, socket layer
15 3402 may be realized as a Winsock™ layer in a Microsoft® Windows® OS
16 environment.

17 In a described implementation, migrator intermediate driver 3414 includes
18 protocol-hardware interfaces 3406 at the junctions to ARP stack 3404(A) and to
19 PNI miniport 3408. Migrator intermediate driver 3414 serves as an offload target
20 in migration offloading procedure 3600. The offload target is a protocol-hardware
21 interface 3406 miniport as illustrated in this example. In a migration uploading
22 procedure 3700 (as in FIG. 37), migrator intermediate driver 3414 serves as an
23 upload diverter.

24 More specifically, migrator intermediate driver 3414 is bound to each
25 physical network interface 3410 through which a TCP connection may be

1 migrated. Migrator intermediate driver 3414 usually operates as a pass-through
2 driver by passing packets upwards or downwards in the network stack without
3 otherwise interacting with the packets. However, migrator intermediate driver
4 3414 does interact with packets related to connection migration (optionally
5 including subsequently tunneled packets).

6 Responsibilities of migrator intermediate driver 3414 include: (i) the
7 acceptance of migrate offload requests; (ii) the aggregation of the protocol state
8 information that is related to the TCP connection being migrated as compiled from
9 the specific protocol stacks 3404(), along with acknowledged data to produce the
10 connection state information; and (iii) the transmission of the aggregated
11 connection state to a targeted device 3500 for a migration uploading procedure
12 3700. A reliable wire protocol for such transmission may be shared with that used
13 by the session tracking components 2002 and 2010 to send and receive session
14 information messages 2008 (e.g., as described above with reference to FIG. 20).

15 Another responsibility of migrator intermediate driver 3414 (e.g., in a
16 migration uploading procedure 3700) is to initiate the uploading of migrated
17 connections that it receives from other devices and to buffer any incoming packets
18 related to the migrating connection while it is in the process of being uploaded. To
19 upload the connection, migrator intermediate driver 3414 sends an upload request
20 to migrator shim 3412. Migrator shim 3412 issues an inject call down into
21 protocol stack 3404 at TCP stack 3404(A) to instantiate the connection in the
22 protocol stack 3404 portion of the network stack.

23 Migrator shim 3412 exposes a transport layer client interface to TCP stack
24 3404(T) and exposes a transport layer provider interface to socket layer 3402.
25 Migrator shim 3412 has two roles: (i) to initiate connection migration offload

1 procedure 3600 on an originating device 3400 and subsequently migration upload
2 procedure 3700 on a targeted device 3500 and (ii) to mediate the classification
3 process between a host application 316 program, a load-balancing classifier 304
4 program, and socket layer 3402. Migrator shim 3412 and migrator intermediate
5 driver 3414 are both further described below with reference to FIGS. 36 and 37.

6 For an exemplary migration offloading procedure 3600, the migration of a
7 TCP connection is performed after classifier 304 classifies the incoming TCP
8 connection using one, two, or more packets thereof. Migration offloading
9 procedure 3600 is described at points <1> through <7>.

10 At <1>, an initialization is performed prior to classification operations.
11 Protocol stack 3404 makes queries at protocol-hardware interface layer 3406 to
12 determine what offloading capabilities, if any, are available. Migrator
13 intermediate driver 3414 indicates that connection migration offloading is
14 available and propagates the query down to PNI miniport 3408. If a TCP chimney
15 offload ability is provided by a physical network interface 3410, PNI miniport
16 3408 also so indicates. TCP chimney offload enables some TCP/IP processing to
17 be offloaded to the hardware of physical network interface 3410 and involves
18 some compiling of protocol state 3420. Consequently, some compiling and
19 aggregation logic may be shared between the two offloading mechanisms.

20 At <2>, once a TCP connection has been classified, classifier 304 initiates a
21 TCP connection migration to a selected host 108. Specifically, a migration
22 command indicating a targeted device 3500 is issued via socket layer 3402 to
23 migrator shim 3412.

24 At <3>, migrator shim 3412 initiates TCP connection migration to compile
25 the TCP protocol state. Specifically, migrator shim 3412 invokes a TCP initiate

1 migrate offload API (or more generally a migrate connection function call or
2 migrate connection command). This routine compiles the relevant state for the
3 specified TCP connection that is used to reinstate the connection on the targeted
4 device 3500. The compiled protocol state 3420 includes state from the
5 intermediate stack layers, including TCP stack 3404(T), IP stack 3404(I), and ARP
6 stack 3404(A).

7 At <4>, once protocol stack 3404 has compiled protocol state 3420 for the
8 TCP connection being migrated, it invokes an initiate migrate offload API on the
9 miniport to which it is bound; in this example, that miniport is migrator
10 intermediate driver 3414. However, in practice, there may be other intermediate
11 drivers inserted between protocol stack 3404 and migrator intermediate driver
12 3414, such as IP QoS. If so, those IM drivers may participate in the migration, if
13 relevant, by compiling/aggregating their state to the connection state information
14 for the connection being migrated. Intermediate drivers continue to propagate the
15 initiate migrate offload call down the network stack, which eventually results in
16 execution of a migrate offload handler at migrator intermediate driver 3414. At
17 this point, migrator intermediate driver 3414 also aggregates any acknowledged
18 data with the remaining connection state for transfer of the TCP connection to
19 targeted device 3500.

20 At <5>, after storing/copying connection state information for the TCP
21 connection being migrated, migrator intermediate driver 3414 notifies the network
22 stack that the migration is in its final stages by invoking an initiate migrate offload
23 complete API. This initiate migrate offload complete API follows the reverse path
24 up the network stack, through the same intermediate drivers (if any), and
25 eventually to protocol stack 3404. As each layer processes this call, state

1 information that is associated with the migrated connection may be released. Until
2 the processing of this call is complete, each layer may send updating notifications
3 down the network stack to update any part of the connection state that has changed
4 since the migration was initiated.

5 At <6>, when the initiate migrate offload complete routine reaches TCP
6 stack 3404(T), TCP silently (i.e., no reset is sent to client 108) closes the
7 connection, flushing all state associated with the migrated connection, and
8 propagates the initiate migrate offload complete call to migrator shim 3412. At
9 this point, the network stack is free of any residual knowledge of the migrated
10 TCP connection.

11 At <7>, when the initiate migrate offload complete call returns to migrator
12 intermediate driver 3414 (via the migrator shim 3412 portion of connection
13 migrator 310), the migration of the TCP connection from originating device 3400
14 to targeted device 3500 may commence with the transfer of the connection state
thereto. The connection state may be transferred asynchronously and reliably.

16 Once migration is initiated, originating device 3400 is also responsible for
17 ensuring that subsequent data from client 108 is forwarded to target device 3500.
18 Consequently, even after the connection is successfully migrated to the target, the
19 originator retains some amount of state for the connection (e.g., a routing table
20 entry) in order to properly route subsequent packets to the target. When the
21 connection is terminated, the target notifies the originator to enable it to purge
22 whatever residual state remains for the migrated connection.

23 Furthermore, as a consequence of the asynchronous nature of the
24 connection migration, data packets for the migrating connection that are forwarded
25 by originating device 3400 (or a forwarder designated thereby if a separate device)

1 may start arriving at targeted device 3500 before targeted device 3500 receives the
2 migrated connection state. Migrator intermediate driver 3414 at targeted device
3 3500 is responsible for buffering those packets until the associated migrated
4 connection is established on targeted device 3500.

5 FIG. 37 illustrates an exemplary approach to an uploading procedure 3700
6 for a connection migration. Migration uploading procedure 3700 illustrates
7 additional exemplary details for a connection migration by targeted device 3500.

8 When a migrated connection arrives at targeted device 3500, it is relayed to
9 migrator intermediate driver 3414 for processing. After amalgamating and
10 assimilating the migrated connection state, migrator intermediate driver 3414, in
11 conjunction with migrator shim 3412, injects the migrated connection into the
12 local network stack in a manner transparent to application 316. For an exemplary
13 migration uploading procedure 3700, the migration of a TCP connection at points
14 <1> through <8> is described.

15 At <1>, as described above with reference to migration offloading
16 procedure 3600, an initialization is performed prior to application hosting
17 operations. Specifically, protocol stack 3404 makes queries regarding what
18 offloading capabilities, if any, are available. Migrator intermediate driver 3414
19 fills in the TCP connection migration support query to indicate that connection
20 migration uploading is available and also propagates the query down to PNI
21 miniport 3408 for possible TCP chimney offload capabilities.

22 At <2>, when connection migration data arrives at target device 3500, the
23 connection migration information (e.g., a bundled binary blob 3422) is delivered
24 to migrator intermediate driver 3414. Migrator intermediate driver 3414 re-
25 assembles the connection state, matches it up with any associated data that has

1 arrived during the migration, and prepares for the upload onto the network stack.
2 Any data from client 102 that arrives during the process of uploading the migrated
3 connection is buffered by migrator intermediate driver 3414. Upon successful
4 completion of the migration, the data will be delivered to application 316.

5 At <3>, to initiate the upload of the migrated connection into the local
6 network stack, migrator intermediate driver 3414 notifies migrator shim 3412 that
7 a migrated connection request has arrived. Migrator intermediate driver 3414 also
8 delivers the connection state (or at least protocol state 3420) to migrator shim
9 3412.

10 At <4>, migrator shim 3412 initiates the upload of the migrated connection
11 by invoking a TCP initiate inject routine (or more generally an infuse protocol
12 state routine) and by providing the migrated protocol state 3420 to TCP stack
13 3404(T). At <5>, TCP/IP recreates the migrated connection throughout protocol
14 stack 3404 using the provided protocol state 3420. This protocol state 3420 may
15 include one or more of transport state (TCP), path state (IP), neighbor and next-
16 hop state (ARP), and so forth.

17 At <6>, if the migrated connection is successfully reestablished on target
18 device 3500, TCP initiates a connect event to a client portion of migrator shim
19 3412 to indicate that a new connection has been established. There are a multitude
20 of possible reasons for failure, but common reasons may include the lack of a
21 corresponding listener, routing failure, etc. In these cases where the network stack
22 is unable to reestablish the migrated connection, no connect event is indicated and
23 a failure status is specified in the initiate inject complete call. Connection
24 migrator 310 is responsible for cleaning up the migration and for sending a reset
25 notification back to client 102 to abandon the connection.

1 At <7>, migrator shim 3412 acts as a provider to propagate the connect
2 event to socket layer 3402 so as to indicate to the listening application 316 that a
3 new connection has been established. If the application 316 accepts the
4 connection, it processes the requests and responds through normal read and write
5 socket operations; application 316 can be unaware that the connection was
6 migrated. If the connection is not accepted by the application 316, TCP terminates
7 the connection but does not send a reset notification back to client 102. Again, a
8 failure status is specified in the initiate inject complete call, and connection
9 migrator 310 is responsible for cleaning up the migration and for sending a reset
10 notification back to client 102 to abandon the connection.

11 A special situation arises when application 316 and classifier 304 are co-
12 located on the same device: migrator shim 3412 may referee between them.
13 When both classes of programs reside on the same host 108, they may both be
14 listening to the same IP address(es) and port(s). However, TCP typically has one
15 listener per unique IP address and port. Consequently, migrator shim 3412 can
16 obscure a configuration where two programs are listening on the same IP address
17 and port by multiplexing the two sockets into a single listener at the TCP layer.

18 In such a case, when connect events arrive at the client portion of migrator
19 shim 3412, migrator shim 3412 as a provider determines on which listening socket
20 to deliver the connect notification at socket layer 3402. If there is only one socket
21 listening to the corresponding IP address and port, then that socket receives the
22 connect event. If there is more than one socket listening, then the recipient
23 depends on the context in which the connect event is indicated. If the connect
24 event is a brand new connection for a virtual IP address, then the connect event is
25 delivered to classifier 304; if the connect event is for a dedicated IP address (non-

1 load-balanced IP address) or the result of uploading a migrated connection, then
2 the connect event is delivered to the target application 316.

3 At <8>, once the injection of the migrated connection is complete, TCP
4 notifies migrator shim 3412 by invoking the provided initiate inject complete
5 handler. A status code is provided to notify migrator shim 3412 whether or not the
6 connection was successfully uploaded. If uploading of the migrated connection
7 fails, connection migrator 310 is responsible for cleaning up the migration and for
8 notifying client 102 that the connection has been abandoned by sending it a reset.
9 If the migrated connection was successfully injected into the local network stack,
10 migrator intermediate driver 3414 may begin delivering any buffered data from
11 client 102 by passing the received packet(s) up through the packet receive path of
12 protocol-hardware interface 3406.

13 When a migrated connection is terminated (because uploading failed,
14 because the migrated connection is subsequently closed through normal means,
15 etc.), target device 3500 notifies originating device 3400. Originating device 3400
16 uses these notifications to more efficiently and reliably clean out lingering state for
17 migrated connections, including routing table entries. Therefore, to account for
18 successfully migrated connections which terminate arbitrarily in the future,
19 migrator shim 3412 may monitor their activity and notify migrator intermediate
20 driver 3414 when the sockets therefor are closed.

21 FIG. 38 illustrates an exemplary approach to packet tunneling between a
22 forwarder 302 and a host 108. Encapsulated packets 3808 may be tunneled from
23 forwarder 302 to host 108 without incurring overhead for each packet transmitted.
24 As described further below, the tunneling is effectuated using a flow identifier
25 3814 and encapsulation mapping tables 3806 and 3810 of tunnelers 312(F) and

1 312(H), respectively, of forwarder 302 and host 108, respectively. Flow identifier
2 3814 is inserted into encapsulated packets 3808.

3 As noted above with reference to FIG. 32, packets for a connection that
4 arrive subsequent to a connection migration may be routed by forwarder 302 to
5 host 108(1) using tunneling by a tunneler 312. At (8) (of FIG. 32), forwarder 302
6 forwards such subsequent packets from forwarder 302 having a network address
7 of “F” to host 108(1) having a network address of “H1”. As described above with
8 reference to FIG. 4, forwarder 302 may perform NAT, half-NAT, tunneling, etc. in
9 order to route the incoming packets to host 108(1).

10 Such incoming packets include a destination IP address of the virtual IP
11 (“VIP”) address and a source IP address of “C1” for packets arriving from client
12 102(1). The packets being routed to host 108(1) have a destination IP address of
13 H1 and a source address of C1 (for half-NAT) or “F” (for full NAT). This re-
14 writing of the addresses can interfere with some protocols that expect both of
15 client 102(1) and host 108(1) to have identical views of the source and destination
16 addresses.

17 Furthermore, at least with respect to full NAT, return paths from host
18 108(1) to client 102(1) that do not run through forwarder 302 are prohibitive
19 because host 108(1) does not know the address of client 102(1). Direct paths from
20 host 108(1) to client 102(1) are desirable in situations in which traffic from host
21 108(1) to client 102(1) is especially high and/or significantly greater than traffic in
22 the opposite direction (e.g., when host 108(1) provides streaming media to client
23 102(1)).

24 Tunneling by tunneler 312 as described herein can provide for identical
25 views with respect to the source and destination addresses (and ports) for clients

1 102 and applications 316 on hosts 108. By way of example and with reference to
2 FIGS. 34 and 35, tunneler 312 in each of forwarder 302 and host 108 may operate
3 as part of or in conjunction with a migrator intermediate driver 3414 of a
4 connection migrator 310.

5 In a described implementation for FIG. 38, connection migrator 310
6 provides an encapsulation mapping 3812 between a flow identifier 3814 and a
7 TCP/IP 4-tuple 3804. Connection migrator 310 may be associated with a classifier
8 304, and connection migrator 310 (optionally along with such a classifier 304)
9 may be located on a same device as forwarder 302. Alternatively, connection
10 migrator 310 (as well as the classifier 304) may be located on a different device
11 from forwarder 302. Encapsulation mapping 3812 may alternatively be provided
12 by or in conjunction with tunneler 312 functionality that is, for example, located at
13 and/or associated with a classifier 304.

14 By being mapped to a TCP/IP 4-tuple 3804 in encapsulation mapping 3812,
15 flow identifier 3814 serves to identify a flow of encapsulated packets 3808 for a
16 particular connection. TCP/IP 4-tuple 3804 includes network addresses (and ports,
17 etc.) for the source and destination for a particular connection in accordance with a
18 TCP/IP protocol, or any similar or analogous protocol. Flow identifier 3814 is 32
19 bits in a described implementation because this allows the flow identifier to be
20 encoded in the source and destination port fields of the TCP segment header in the
21 tunneled packet, which enables the tunneled packet to be transmitted without any
22 tunneling space overhead. At the destination, the TCP/IP 4-tuple can be
23 determined by looking up the 4-tuple that is linked to the flow identifier as
24 extracted from the source and destination port fields. However, flow identifiers
25

1 3814 of other lengths may alternatively be used, especially for other protocols
2 such as internet RTP, etc.

3 Each flow identifier 3814 can identify a unique connection from the device
4 that is originating the tunneling (which is forwarder 302 in this example). Flow
5 identifiers 3814 may be generated using any appropriate mechanism, such as an
6 incrementing connection counter. Alternatively, the TCP/IP receiver Initial
7 Sequence Number (ISN) generated by the connection migrator can serve as flow
8 identifiers 3814. Furthermore, TCP/IP 4-tuple 3804 is more generally a
9 source/destination pair. Each source value and destination value of an individual
10 source/destination pair may include a network node identifier (e.g., network
11 address, port, some combination thereof, etc.) for the source and destination,
12 respectively, of a given packet propagating on a particular connection.

13 Connection migrator 310 provides encapsulation mapping 3812 to host 108.
14 Tunneler 312(H) at host 108 stores encapsulation mapping 3812 in encapsulation
15 mapping table 3810 as encapsulation mapping entry 3810(1). Tunneler 312(H)
16 can thereafter use flow identifier 3814 to map to and identify the particular
17 connection corresponding to TCP/IP 4-tuple 3804. Encapsulation mapping 3812
18 may optionally be provided to host 108 as part of a bundled binary blob 3422 in a
19 connection migration operation.

20 Forwarder 302 also includes a tunneler 312(F) component with an
21 encapsulation mapping table 3806. Encapsulation mapping table 3806 stores an
22 encapsulation mapping entry 3806(1) that links/maps TCP/IP 4-tuple 3804 for a
23 particular connection to a flow identifier 3814. Tunneler 312(F) also receives the
24 mapping information for encapsulation mapping entry 3806(1) from connection
25 migrator 310 (e.g., as an encapsulation mapping 3812).

1 Although only one encapsulation mapping entry 3806(1) and 3810(1) is
2 shown, each of encapsulation mapping table 3806 and encapsulation mapping
3 table 3810 may have multiple such entries. These encapsulation mapping tables
4 3806 and 3810 may be combined with other information, such as tables for session
5 information of session tracker 308.

6 When a transmitting device (such as forwarder 302) and a receiving device
7 (such as host 108) of encapsulated packets 3808 only tunnel between each other,
8 the encapsulation mapping tables thereof likely have the same encapsulation
9 mapping entries. Otherwise, encapsulation mapping table 3806 and encapsulation
10 mapping table 3810 likely have a different total set of encapsulation mapping
11 entries 3806() and encapsulation mapping entries 3810(), respectively.

12 In operation, an incoming packet 3802 for a particular connection is
13 received at forwarder 302. The particular connection is associated with TCP/IP 4-
14 tuple 3804. Incoming packet 3802 includes TCP/IP 4-tuple 3804 with a source IP
15 address (of a client 102), a destination IP address (the virtual IP), a source TCP
16 port (of the client 102), and a destination TCP port.

17 Tunneler 312(F) accepts incoming packet 3802 for tunneling to host 108.
18 Using TCP/IP 4-tuple 3804, tunneler 312(F) accesses encapsulation mapping table
19 3806 to locate encapsulation mapping entry 3806(1). Flow identifier 3814 is
20 extracted from encapsulation mapping entry 3806(1) as being linked/mapped to
21 TCP/IP 4-tuple 3804.

22 To create encapsulated packet 3808, tunneler 312(F) inserts flow identifier
23 3814 into the source and destination port portions of the TCP/IP 4-tuple header.
24 These two TCP portions are 16 bits each, which allows a 32-bit flow identifier
25 3814 to be inserted. Also, for the source IP address portion of the TCP/IP 4-tuple

1 header, tunneler 312(F) inserts the IP address “F” of forwarder 302. For the
2 destination IP address portion of the TCP/IP 4-tuple header, tunneler 312(F) inserts
3 the IP address “H” of host 108.

4 Forwarder 302 routes/transmits encapsulated packet 3808 to host 108, and
5 host 108 receives encapsulated packet 3808 from forwarder 302. The tunneler
6 312(H) component at host 108 detects that encapsulated packet 3808 is a tunneled
7 packet that is to be de-encapsulated.

8 Flow identifier 3814 is extracted from encapsulated packet 3808 and used
9 to look up the corresponding TCP/IP 4-tuple 3804 that is linked thereto in
10 encapsulation mapping entry 3810(1) of encapsulation mapping table 3810.
11 TCP/IP 4-tuple 3804 is used by tunneler 312(H) to recreate the TCP/IP 4-tuple
12 3804 header as originally received in incoming packet 3802 at forwarder 302.

13 Specifically, the IP address F of forwarder 302 is replaced with the source
14 IP address, and the IP address H of host 108 is replaced with the destination IP
15 address. Furthermore, flow identifier 3814 is replaced by the source TCP port and
16 the destination TCP port. The de-encapsulated packet is then indicated up the
17 network stack of host 108 to the targeted application 316.

18 More generally, a portion of a packet header, including a portion of a
19 source/destination pair, for a given packet that is not necessarily used for
20 communicating the given packet may be used to carry a flow identifier 3814. By
21 pre-providing at least part of the source/destination pair at host 108, a flow
22 identifier 3814 may be employed to tunnel (e.g., encapsulate and/or de-
23 encapsulate) packets without incurring an encapsulation overhead on each packet.
24 Furthermore, packets that are full-size with respect to a given protocol may be
25 tunneled without being fragmented.

1 FIG. 39 is a flow diagram 3900 that illustrates an exemplary method for
2 packet tunneling between a first device and a second device. For example, the
3 first device and the second device may correspond to an originating device 3400
4 and a target device 3500, respectively, of load balancing infrastructure 106 and a
5 cluster of hosts 108, respectively. Nevertheless, tunneling may be employed in
6 non-load-balancing implementations.

7 Flow diagram 3900 includes twelve blocks 3902-3924. Although the
8 actions of flow diagram 3900 may be performed in other environments and with a
9 variety of software schemes, FIGS. 1-3, 32, 34, 35, and 38 are used in particular to
10 illustrate certain aspects and examples of the method.

11 At block 3902, a mapping of a flow identifier-to-TCP/IP 4-tuple is sent to a
12 target device from an originating device. For example, originating device 3400
13 may send an encapsulation mapping 3812 that links a flow identifier 3814 to a
14 TCP/IP 4-tuple 3804. At block 3914, the mapping of the flow identifier-to-the
15 TCP/IP 4-tuple is received at the target device from the originating device. For
16 example, target device 3500 receives encapsulation mapping 3812 that links flow
17 identifier 3814 to TCP/IP 4-tuple 3804 from originating device 3400.

18 Alternatively, target device 3500 may receive encapsulation mapping 3812
19 from another device. As indicated by dashed arrows 3926 and 3928, the actions of
20 blocks 3904-3912 and blocks 3916-3924 can occur at some time after the actions
21 of blocks 3902 and 3914, respectively.

22 At block 3904, an incoming packet is received at the originating device
23 from a client. For example, an incoming packet 3802 having a header with
24 TCP/IP 4-tuple 3804 may be received at originating device 3400 from a client 102.
25 At block 3906, a flow identifier is looked up for a connection corresponding to the

1 client's packet using the TPC/IP 4-tuple of the incoming packet. For example,
2 flow identifier 3814 may be looked up for the connection with client 102 using
3 TCP/IP 4-tuple 3804 that is mapped thereto in an encapsulation mapping entry
4 3806(1) of an encapsulation mapping table 3806.

5 At block 3908, the source IP and destination IP of the incoming packet are
6 replaced with an originating IP address of the originating device and a target IP
7 address of the target device, respectively. For example, originating device 3400
8 may replace the IP address portions of the TCP/IP 4-tuple 3804 portion of a header
9 of incoming packet 3802 with IP addresses of originating device 3400 and target
10 device 3500.

11 At block 3910, the source port and the destination port of the incoming
12 packet are replaced with the flow identifier. For example, originating device 3400
13 may replace source and destination TCP ports of the TCP/IP 4-tuple 3804 portion
14 of the header of incoming packet 3802 with flow identifier 3814. At block 3912,
15 the encapsulated packet is sent from the originating device to the target device.
16 For example, originating device 3400 may send an encapsulated packet 3808 to
17 target device 3500.

18 At block 3916, the encapsulated packet is received at the target device from
19 the originating device. For example, target device 3500 may receive the
20 encapsulated packet 3808 from originating device 3400. At block 3918, the
21 TCP/IP 4-tuple is looked up for the connection corresponding to the packet
22 received from the client using the flow identifier. For example, target device 3500
23 may access an encapsulation mapping table 3810 at an encapsulation mapping
24 entry 3810(1) that maps flow identifier 3814 to TCP/IP 4-tuple 3804.

25

1 At block 3920, the originating IP address and the target IP address are
2 replaced with the source IP address and the destination IP address, respectively,
3 using the looked-up TCP/IP 4-tuple. For example, target device 3500 may replace
4 the IP addresses of originating device 3400 and target device 3500 in encapsulated
5 packet 3808 with the source IP address and the destination IP address from TCP/IP
6 4-tuple 3804 as attained from encapsulation mapping table 3810.

7 At block 3922, the flow identifier is replaced with the source port and the
8 destination port of the incoming packet using the looked up TCP/IP 4-tuple. For
9 example, target device 3500 may replace flow identifier 3814 in encapsulated
10 packet 3808 with the source TCP port and the destination TCP port from TCP/IP
11 4-tuple 3804. At block 3924, the client's packet is indicated up to an application at
12 the target device. For example, a de-encapsulated version of encapsulated packet
13 3808, or incoming packet 3802, is indicated up to application 316 of target device
14 3500.

15 The actions, aspects, features, components, etc. of FIGS. 1-39 are
16 illustrated in diagrams that are divided into multiple blocks. However, the order,
17 interconnections, layout, etc. in which FIGS. 1-39 are described and/or shown is
18 not intended to be construed as a limitation, and any number of the blocks can be
19 combined, rearranged, augmented, omitted, etc. in any manner to implement one
20 or more systems, methods, devices, procedures, media, APIs, apparatuses,
21 arrangements, etc. for network load balancing. Furthermore, although the
22 description herein includes references to specific implementations (and the
23 exemplary operating environment of FIG. 40), the illustrated and/or described
24 implementations can be implemented in any suitable hardware, software,
25 firmware, or combination thereof and using any suitable network organization(s),

1 transport/communication protocols(s), application programming interface(s)
2 (APIs), client-server architecture(s), and so forth.

3 **Exemplary Operating Environment for Computer or Other Device**

4 FIG. 40 illustrates an exemplary computing (or general device) operating
5 environment 4000 that is capable of (fully or partially) implementing at least one
6 system, device, apparatus, component, arrangement, protocol, approach, method,
7 procedure, media, API, some combination thereof, etc. for network load balancing
8 as described herein. Operating environment 4000 may be utilized in the computer
9 and network architectures described below or in a stand-alone situation.

10 Exemplary operating environment 4000 is only one example of an
11 environment and is not intended to suggest any limitation as to the scope of use or
12 functionality of the applicable device (including computer, network node,
13 entertainment device, mobile appliance, general electronic device, etc.)
14 architectures. Neither should operating environment 4000 (or the devices thereof)
15 be interpreted as having any dependency or requirement relating to any one or to
16 any combination of components as illustrated in FIG. 40.

17 Additionally, network load balancing may be implemented with numerous
18 other general purpose or special purpose device (including computing system)
19 environments or configurations. Examples of well known devices, systems,
20 environments, and/or configurations that may be suitable for use include, but are
21 not limited to, personal computers, server computers, thin clients, thick clients,
22 personal digital assistants (PDAs) or mobile telephones, watches, hand-held or
23 laptop devices, multiprocessor systems, microprocessor-based systems, set-top
24 boxes, programmable consumer electronics, video game machines, game consoles,
25 portable or handheld gaming units, network PCs, minicomputers, mainframe

1 computers, network nodes, distributed or multi-processing computing
2 environments that include any of the above systems or devices, some combination
3 thereof, and so forth.

4 Implementations for network load balancing may be described in the
5 general context of processor-executable instructions. Generally, processor-
6 executable instructions include routines, programs, protocols, objects, interfaces,
7 components, data structures, etc. that perform and/or enable particular tasks and/or
8 implement particular abstract data types. Network load balancing, as described in
9 certain implementations herein, may also be practiced in distributed processing
10 environments where tasks are performed by remotely-linked processing devices
11 that are connected through a communications link and/or network. Especially in a
12 distributed computing environment, processor-executable instructions may be
13 located in separate storage media, executed by different processors, and/or
14 propagated over transmission media.

15 Exemplary operating environment 4000 includes a general-purpose
16 computing device in the form of a computer 4002, which may comprise any (e.g.,
17 electronic) device with computing/processing capabilities. The components of
18 computer 4002 may include, but are not limited to, one or more processors or
19 processing units 4004, a system memory 4006, and a system bus 4008 that couples
20 various system components including processor 4004 to system memory 4006.

21 Processors 4004 are not limited by the materials from which they are
22 formed or the processing mechanisms employed therein. For example, processors
23 4004 may be comprised of semiconductor(s) and/or transistors (e.g., electronic
24 integrated circuits (ICs)). In such a context, processor-executable instructions may
25 be electronically-executable instructions. Alternatively, the mechanisms of or for

1 processors 4004, and thus of or for computer 4002, may include, but are not
2 limited to, quantum computing, optical computing, mechanical computing (e.g.,
3 using nanotechnology), and so forth.

4 System bus 4008 represents one or more of any of many types of wired or
5 wireless bus structures, including a memory bus or memory controller, a point-to-
6 point connection, a switching fabric, a peripheral bus, an accelerated graphics port,
7 and a processor or local bus using any of a variety of bus architectures. By way of
8 example, such architectures may include an Industry Standard Architecture (ISA)
9 bus, a Micro Channel Architecture (MCA) bus, an Enhanced ISA (EISA) bus, a
10 Video Electronics Standards Association (VESA) local bus, a Peripheral
11 Component Interconnects (PCI) bus also known as a Mezzanine bus, some
12 combination thereof, and so forth.

13 Computer 4002 typically includes a variety of processor-accessible media.
14 Such media may be any available media that is accessible by computer 4002 or
15 another (e.g., electronic) device, and it includes both volatile and non-volatile
16 media, removable and non-removable media, and storage and transmission media.

17 System memory 4006 includes processor-accessible storage media in the
18 form of volatile memory, such as random access memory (RAM) 4040, and/or
19 non-volatile memory, such as read only memory (ROM) 4012. A basic
20 input/output system (BIOS) 4014, containing the basic routines that help to
21 transfer information between elements within computer 4002, such as during start-
22 up, is typically stored in ROM 4012. RAM 4010 typically contains data and/or
23 program modules/instructions that are immediately accessible to and/or being
24 presently operated on by processing unit 4004.

25

1 Computer 4002 may also include other removable/non-removable and/or
2 volatile/non-volatile storage media. By way of example, FIG. 40 illustrates a hard
3 disk drive or disk drive array 4016 for reading from and writing to a (typically)
4 non-removable, non-volatile magnetic media (not separately shown); a magnetic
5 disk drive 4018 for reading from and writing to a (typically) removable, non-
6 volatile magnetic disk 4020 (e.g., a “floppy disk”); and an optical disk drive 4022
7 for reading from and/or writing to a (typically) removable, non-volatile optical
8 disk 4024 such as a CD, DVD, or other optical media. Hard disk drive 4016,
9 magnetic disk drive 4018, and optical disk drive 4022 are each connected to
10 system bus 4008 by one or more storage media interfaces 4026. Alternatively,
11 hard disk drive 4016, magnetic disk drive 4018, and optical disk drive 4022 may
12 be connected to system bus 4008 by one or more other separate or combined
13 interfaces (not shown).

14 The disk drives and their associated processor-accessible media provide
15 non-volatile storage of processor-executable instructions, such as data structures,
16 program modules, and other data for computer 4002. Although exemplary
17 computer 4002 illustrates a hard disk 4016, a removable magnetic disk 4020, and a
18 removable optical disk 4024, it is to be appreciated that other types of processor-
19 accessible media may store instructions that are accessible by a device, such as
20 magnetic cassettes or other magnetic storage devices, flash memory, compact
21 disks (CDs), digital versatile disks (DVDs) or other optical storage, RAM, ROM,
22 electrically-erasable programmable read-only memories (EEPROM), and so forth.
23 Such media may also include so-called special purpose or hard-wired IC chips. In
24 other words, any processor-accessible media may be utilized to realize the storage
25 media of the exemplary operating environment 4000.

1 Any number of program modules (or other units or sets of
2 instructions/code) may be stored on hard disk 4016, magnetic disk 4020, optical
3 disk 4024, ROM 4012, and/or RAM 4040, including by way of general example,
4 an operating system 4028, one or more application programs 4030, other program
5 modules 4032, and program data 4034.

6 A user may enter commands and/or information into computer 4002 via
7 input devices such as a keyboard 4036 and a pointing device 4038 (e.g., a
8 "mouse"). Other input devices 4040 (not shown specifically) may include a
9 microphone, joystick, game pad, satellite dish, serial port, scanner, and/or the like.
10 These and other input devices are connected to processing unit 4004 via
11 input/output interfaces 4042 that are coupled to system bus 4008. However, input
12 devices and/or output devices may instead be connected by other interface and bus
13 structures, such as a parallel port, a game port, a universal serial bus (USB) port,
14 an infrared port, an IEEE 1394 ("Firewire") interface, an IEEE 802.11 wireless
15 interface, a Bluetooth® wireless interface, and so forth.

16 A monitor/view screen 4044 or other type of display device may also be
17 connected to system bus 4008 via an interface, such as a video adapter 4046.
18 Video adapter 4046 (or another component) may be or may include a graphics
19 card for processing graphics-intensive calculations and for handling demanding
20 display requirements. Typically, a graphics card includes a graphics processing
21 unit (GPU), video RAM (VRAM), etc. to facilitate the expeditious display of
22 graphics and performance of graphics operations. In addition to monitor 4044,
23 other output peripheral devices may include components such as speakers (not
24 shown) and a printer 4048, which may be connected to computer 4002 via
25 input/output interfaces 4042.

1 Computer 4002 may operate in a networked environment using logical
2 connections to one or more remote computers, such as a remote computing device
3 4050. By way of example, remote computing device 4050 may be a personal
4 computer, a portable computer (e.g., laptop computer, tablet computer, PDA,
5 mobile station, etc.), a palm or pocket-sized computer, a watch, a gaming device, a
6 server, a router, a network computer, a peer device, another network node, or
7 another device type as listed above, and so forth. However, remote computing
8 device 4050 is illustrated as a portable computer that may include many or all of
9 the elements and features described herein with respect to computer 4002.

10 Logical connections between computer 4002 and remote computer 4050 are
11 depicted as a local area network (LAN) 4052 and a general wide area network
12 (WAN) 4054. Such networking environments are commonplace in offices,
13 enterprise-wide computer networks, intranets, the Internet, fixed and mobile
14 telephone networks, ad-hoc and infrastructure wireless networks, other wireless
15 networks, gaming networks, some combination thereof, and so forth. Such
16 networks and communications connections are examples of transmission media.

17 When implemented in a LAN networking environment, computer 4002 is
18 usually connected to LAN 4052 via a network interface or adapter 4056. When
19 implemented in a WAN networking environment, computer 4002 typically
20 includes a modem 4058 or other means for establishing communications over
21 WAN 4054. Modem 4058, which may be internal or external to computer 4002,
22 may be connected to system bus 4008 via input/output interfaces 4042 or any
23 other appropriate mechanism(s). It is to be appreciated that the illustrated network
24 connections are exemplary and that other means of establishing communication
25 link(s) between computers 4002 and 4050 may be employed.

1 Furthermore, other hardware that is specifically designed for servers may
2 be employed. For example, SSL acceleration cards can be used to offload SSL
3 computations. Additionally, especially in a network load balancing operating
4 environment, TCP offload hardware and/or packet classifiers on network
5 interfaces or adapters 4056 (e.g., on network interface cards) may be installed and
6 used at server devices.

7 In a networked environment, such as that illustrated with operating
8 environment 4000, program modules or other instructions that are depicted
9 relative to computer 4002, or portions thereof, may be fully or partially stored in a
10 remote media storage device. By way of example, remote application programs
11 4060 reside on a memory component of remote computer 4050 but may be usable
12 or otherwise accessible via computer 4002. Also, for purposes of illustration,
13 application programs 4030 and other processor-executable instructions such as
14 operating system 4028 are illustrated herein as discrete blocks, but it is recognized
15 that such programs, components, and other instructions reside at various times in
16 different storage components of computing device 4002 (and/or remote computing
17 device 4050) and are executed by processor(s) 4004 of computer 4002 (and/or
18 those of remote computing device 4050).

19 Although systems, media, devices, methods, procedures, apparatuses,
20 techniques, schemes, approaches, procedures, arrangements, and other
21 implementations have been described in language specific to structural, logical,
22 algorithmic, and functional features and/or diagrams, it is to be understood that the
23 invention defined in the appended claims is not necessarily limited to the specific
24 features or diagrams described. Rather, the specific features and diagrams are
25 disclosed as exemplary forms of implementing the claimed invention.